

IBM Spectrum Protect Knowledge Center Version 8.1.5



Spis treści

Witamy	1
Dostępność	1
Pakiety produktów i produkty pokrewne	2
Pliki PDF	5
Aktualizacje w tej wersji	5
IBM Spectrum Protect concepts	5
IBM Spectrum Protect overview	6
Data protection components	6
Data protection services	7
Data protection management processes	9
User interfaces	11
Data storage concepts	12
Data storage devices	13
Storage pools	15
Data transport to storage	19
Data protection strategies	21
Backup storage space minimization	22
Disaster protection strategies	23
Disaster recovery concepts	26
Data protection solutions	28
Selecting a data protection solution	28
Single-site disk solution	28
Multisite disk solution	29
Multisite appliance solution	30
Tape solution	31
Solutions comparison	32
Solution roadmap	34
Single-site disk solution	34
Planning	35
Selecting a system size	36
System requirements for a single-site disk solution	36
Hardware requirements	36
Software requirements	38
Planning worksheets	39
Planning for storage	47
Planning for security	47
Planning for administrator roles	48
Planning for secure communications	48
Planning for storage of encrypted data	49
Planning firewall access	49
Implementation	50
Setting up the system	50
Configuring the storage hardware	51
Installing the server operating system	51

Installing on AIX systems	51
Installing on Linux systems	53
Installing on Windows systems	56
Configuring multipath I/O	56
AIX systems	57
Linux systems	57
Windows systems	58
Creating the user ID for the server	59
Preparing file systems for the server	60
AIX systems	60
Linux systems	61
Windows systems	62
Installing the server and Operations Center	62
Installing on AIX and Linux systems	63
Installing on Windows systems	63
Configuring the server and the Operations Center	64
Configuring the server instance	65
Installing the backup-archive client	65
Setting options for the server	66
Configuring secure communications with Transport Layer Security	67
Configuring the Operations Center	67
Registering the product license	68
Configuring data deduplication	68
Defining data retention rules for your business	69
Defining schedules for server maintenance activities	69
Defining client schedules	71
Installing and configuring backup-archive clients	71
Registering and assigning clients to schedules	72
Installing the client management service	72
Verifying that the client management service is installed correctly	73
Configuring the Operations Center to use the client management service	74
Completing the implementation	74
Monitoring	75
Daily checklist	75
Periodic checklist	82
Verifying license compliance	87
Tracking system status by using email reports	88
Managing	89
Managing the Operations Center	89
Adding and removing spoke servers	90
Adding a spoke server	90
Removing a spoke server	90
Starting and stopping the web server	91
Restarting the initial configuration wizard	91
Changing the hub server	92
Restoring the configuration to the preconfiguration state	92
Protecting applications, virtual machines, and systems	94
Adding clients	94
Selecting the client software and planning the installation	95
Specifying rules for backing up and archiving client data	96
Viewing policies	97
Editing policies	97
Scheduling backup and archive operations	98
Registering clients	99
Installing and configuring clients	100
Configuring the client to run scheduled operations	101

Configuring communications through a firewall	102
Managing client operations	103
Evaluating errors in client error logs	103
Stopping and restarting the client acceptor	104
Resetting passwords	105
Modifying the scope of a client backup	106
Managing client upgrades	106
Decommissioning a client node	107
Deactivating data to free storage space	109
Managing data storage	109
Auditing a storage pool container	109
Managing inventory capacity	110
Managing memory and processor usage	112
Tuning scheduled activities	112
Securing the server	112
Security concepts	113
Managing administrators	115
Changing password requirements	115
Securing the server on the system	116
Restricting user access to the server	117
Limiting access through port restrictions	117
Stopping and starting the server	118
Stopping the server	118
Starting the server for maintenance or reconfiguration tasks	119
Planning to upgrade the server	120
Preparing for an outage	120
Implementing a disaster recovery plan	121
Recovering from system outages	121
Multisite disk solution	122
Planning	122
Selecting a system size	123
Planning the sites	124
System requirements for a multisite disk solution	125
Hardware requirements	125
Software requirements	127
Planning worksheets	128
Planning for storage	136
Planning for security	136
Planning for administrator roles	137
Planning for secure communications	137
Planning for storage of encrypted data	138
Planning firewall access	138
Implementation	139
Setting up the system	140
Configuring the storage hardware	140
Installing the server operating system	140
Installing on AIX systems	140
Installing on Linux systems	142
Installing on Windows systems	145
Configuring multipath I/O	146
AIX systems	146
Linux systems	147
Windows systems	148
Creating the user ID for the server	148
Preparing file systems for the server	149
AIX systems	149

Linux systems	150
Windows systems	151
Installing the server and Operations Center	152
Installing on AIX and Linux systems	152
Installing on Windows systems	153
Configuring the server and the Operations Center	153
Configuring the server instance	154
Installing the backup-archive client	155
Setting options for the server	155
Configuring secure communications with Transport Layer Security	156
Configuring the Operations Center	156
Registering the product license	157
Configuring data deduplication	158
Defining data retention rules for your business	158
Defining schedules for server maintenance activities	158
Defining client schedules	161
Installing and configuring backup-archive clients	161
Registering and assigning clients to schedules	161
Installing the client management service	162
Verifying that the client management service is installed correctly	162
Configuring the Operations Center to use the client management service	163
Configuring the second server	164
Configuring SSL communications between the hub server and a spoke server	164
Adding the second server as a spoke	166
Enabling replication	166
Completing the implementation	166
Monitoring	167
Daily checklist	167
Periodic checklist	174
Verifying license compliance	179
Tracking system status by using email reports	180
Managing	181
Managing the Operations Center	181
Adding and removing spoke servers	182
Adding a spoke server	182
Removing a spoke server	182
Starting and stopping the web server	183
Restarting the initial configuration wizard	184
Changing the hub server	184
Restoring the configuration to the preconfiguration state	184
Protecting applications, virtual machines, and systems	186
Adding clients	186
Selecting the client software and planning the installation	187
Specifying rules for backing up and archiving client data	188
Viewing policies	189
Editing policies	189
Scheduling backup and archive operations	190
Registering clients	191
Installing and configuring clients	192
Configuring the client to run scheduled operations	193
Configuring communications through a firewall	194
Managing client operations	195
Evaluating errors in client error logs	195
Stopping and restarting the client acceptor	196
Resetting passwords	197
Modifying the scope of a client backup	198

Managing client upgrades	198
Decommissioning a client node	199
Deactivating data to free storage space	201
Managing data storage	201
Auditing a storage pool container	202
Managing inventory capacity	202
Managing memory and processor usage	204
Tuning scheduled activities	204
Managing replication	205
Replication compatibility	205
Enabling node replication	206
Protecting data in directory-container storage pools	206
Modifying replication settings	208
Setting different retention policies	208
Securing the server	209
Security concepts	209
Managing administrators	211
Changing password requirements	212
Securing IBM Spectrum Protect on the system	213
Restricting user access to the server	213
Limiting access through port restrictions	214
Stopping and starting the server	214
Stopping the server	215
Starting the server for maintenance or reconfiguration tasks	216
Planning to upgrade the server	216
Preparing for an outage	217
Implementing a disaster recovery plan	217
Recovering from data loss or system outages	217
Restoring the database	220
Recovering damaged data	221
Repairing storage pools	222
Tape solution	222
Planning	223
Tape planning requirements	223
System requirements for a tape-based solution	224
Hardware requirements	224
Software requirements	227
Planning worksheets	228
Planning for disk storage	231
Planning for tape storage	232
Supported tape devices and libraries	232
Supported tape device configurations	233
Data movement between storage devices	233
Library sharing	234
LAN-free data movement	234
Mixed device types in libraries	235
Different media generations in a library	236
Mixed media and storage pools	236
Required definitions for tape storage devices	237
Planning the storage pool hierarchy	237
Offsite data storage	239
Planning for security	240
Planning for administrator roles	240
Planning for secure communications	240
Planning for storage of encrypted data	241
Planning firewall access	241

Implementing	242
Setting up the system	243
Configuring the storage hardware	244
Installing the server operating system	244
Installing on AIX systems	244
Installing on Linux systems	246
Installing on Windows systems	249
Configuring multipath I/O	249
AIX systems	250
Linux systems	251
Windows systems	252
Creating the user ID for the server	252
Preparing file systems for the server	253
AIX systems	253
Linux systems	254
Windows systems	255
Installing the server and Operations Center	256
Installing on AIX and Linux systems	256
Installing on Windows systems	257
Configuring the server and the Operations Center	257
Configuring the server instance	258
Installing the backup-archive client	258
Setting options for the server	259
Security concepts	260
Configuring the Operations Center	262
Registering the product license	262
Defining data retention rules for your business	263
Defining schedules for server maintenance activities	263
Defining client schedules	267
Attaching tape devices for the server	268
Attaching an automated library device to your system	268
Selecting a tape device driver	269
IBM tape device drivers	269
IBM Spectrum Protect tape device drivers	269
Special file names for tape devices	270
Installing and configuring tape device drivers	271
Installing and configuring IBM device drivers for IBM tape devices	271
AIX systems	272
SCSI and Fibre Channel devices	273
Configuring IBM Spectrum Protect device drivers for autochangers	274
Configuring IBM Spectrum Protect device drivers for tape drives	274
Configuring Fibre Channel SAN-attached devices	275
Linux systems	275
Configuring IBM Spectrum Protect passthru drivers for tape devices and libraries	275
Installing zSeries Linux Fibre Channel adapter (zfcp) device drivers	276
Information about your system's SCSI devices	276
Preventing tape labels from being overwritten	277
Windows systems	278
Preparing to use the IBM Spectrum Protect passthru driver for tape devices and libraries	278
Configuring the IBM Spectrum Protect SCSI driver for tape devices and libraries	278
Configuring libraries for use by a server	279
Defining tape devices	280
Defining libraries and drives	280
Defining libraries	281
Defining drives	281
Defining tape device classes	282

Defining LTO device classes	283
Mixing LTO drives and media in a library	283
Mount limits in LTO mixed-media environments	284
Enabling and disabling drive encryption for LTO Generation 4 or later tape drives	285
Defining 3592 device classes	286
Mixing generations of 3592 drives and media in a single library	286
Controlling data-access speeds for 3592 volumes	287
Enabling and disabling 3592 Generation 2 and later drive encryption	288
Configuring library sharing	288
Example: Library sharing for AIX and Linux servers	289
Example: Library sharing for Windows servers	290
Setting up the library manager server	291
Setting up the library client servers	292
Setting up a storage pool hierarchy	293
Protecting applications and systems	294
Configuring LAN-free data movement	294
Encryption methods	295
Controlling tape storage operations	296
How IBM Spectrum Protect fills volumes	297
Specifying the estimated capacity of tape volumes	297
Specifying recording formats for tape media	298
Associating library objects with device classes	298
Controlling media-mount operations for tape devices	298
Controlling the number of simultaneously mounted volumes	299
Controlling the amount of time that a volume remains mounted	300
Controlling the amount of time that the server waits for a drive	300
Preempting operations	300
Mount point preemption	301
Volume access preemption	301
Impacts of device changes on the SAN	302
Displaying device information	302
Write-once, read-many tape media	303
WORM-capable drives	303
Check-in of WORM media	304
Restrictions on WORM media	304
Mount failures with WORM media	304
Relabeling WORM media	304
Removing private WORM volumes from a library	305
Creation of DLT WORM volumes	305
Support for short and normal 3592 WORM tapes	305
Querying a device class for the WORM-parameter setting	305
Troubleshooting problems with devices	305
Completing the implementation	306
Monitoring	306
Daily checklist	307
Periodic checklist	313
Monitoring tape alert messages for hardware errors	318
Preventing errors caused by media incompatibility	319
Operations with cleaner cartridges	319
Verifying license compliance	319
Tracking system status by using email reports	321
Managing	321
Managing the Operations Center	322
Managing client operations	322
Evaluating errors in client error logs	323
Stopping and restarting the client acceptor	323

Resetting passwords	324
Managing client upgrades	325
Decommissioning a client node	326
Deactivating data to free storage space	328
Managing data storage	328
Managing inventory capacity	328
Tuning scheduled activities	330
Optimizing operations by enabling collocation of client files	330
Effects of collocation on operations	332
Selecting volumes with collocation enabled	333
Selecting volumes with collocation disabled	334
Collocation settings	335
Collocation of copy storage pools	335
Planning for and enabling collocation	336
Managing tape devices	337
Preparing removable media	337
Labeling tape volumes	338
Checking storage volumes into a library	338
Checking a single volume into a SCSI library	339
Checking in volumes from library storage slots	340
Checking in storage volumes from library entry/exit ports	340
Checking in volumes by using library bar code readers	341
Checking in volumes	341
Checking volumes into a full library with swapping	341
Private volumes and scratch volumes	342
Element addresses for library storage slots	342
Managing volume inventory	343
Controlling access to volumes	343
Reusing tapes	343
Maintaining a supply of scratch volumes	345
Maintaining a supply of volumes in a library that contains WORM media	345
Manage the volume inventory in automated libraries	346
Changing the status of a volume in an automated library	347
Removing volumes from an automated library	347
Maintaining a supply of scratch volumes in an automated library	347
Managing an overflow location	348
Auditing the volume inventory	349
Partially written volumes	349
Shared library operations	349
Server requests for volumes	350
Managing tape drives	352
Updating drives	352
Data validation during read/write operations to tape	353
Supported drives	354
Enabling and disabling logical block protection	354
Read/write operations to volumes	355
Storage pool management in a tape library	356
Cleaning tape drives	356
Methods for cleaning tape drives	357
Configuring the server for drive cleaning in an automated library	357
Checking a cleaner cartridge into a library	358
Operations with cleaner cartridges	319
Resolving errors that are related to drive cleaning	359
Tape drive replacement	359
Deleting tape drives	360
Replacing drives with others of the same type	360

Migrating data to upgraded drives	361
Securing the server	361
Managing administrators	362
Changing password requirements	362
Securing the server on the system	363
Stopping and starting the server	363
Stopping the server	364
Starting the server for maintenance or reconfiguration tasks	365
Planning to upgrade the server	365
Preparing for an outage	366
Preparing for and recovering from a disaster by using DRM	366
Disaster recovery plan file	367
Recovering the server and client data	369
Recovery drills	370
Restoring the database	371
PDF files	372

Servers

Servers	372
What's new	372
Operations Center updates	374
Server updates	374
Lower the cost of cloud-container storage pools by reclaiming space	375
Manage your storage environment to help you support General Data Protection Regulation compliance strategies	375
Generate data deduplication statistics for specified nodes and file spaces	375
Schedule audit operations to identify corrupted files in a storage pool	376
V8.1 release notes	376
Servers	376
Operations Center	377
Devices	379
V8.1 readme files for fix packs	380
Installing and upgrading	380
Implementing an IBM Spectrum Protect solution	381
Availability of features by operating system	381
Installing and upgrading the server	382
AIX: Installing the server	383
AIX: Planning to install the IBM Spectrum Protect server	383
AIX: What you should know first	384
AIX: Planning for optimal performance	384
AIX: Planning server hardware and operating system	384
AIX: Planning server database disks	387
AIX: Planning server recovery log disks	389
AIX: Planning container storage pools	390
AIX: Planning DISK or FILE storage pools	396
AIX: Planning storage technology	397
AIX: Installation best practices	399
AIX: Minimum system requirements for AIX systems	400
AIX: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system	400
AIX: IBM Installation Manager	401
AIX: Worksheets for planning details for the server	402
AIX: Capacity planning	402
AIX: Database space requirements	403
AIX: Maximum number of files	403
AIX: Storage pool capacity	405
AIX: The database manager and temporary space	405
AIX: Recovery log space requirements	406

AIX: Active and archive log space	406
AIX: Example: Basic client-store operations	407
AIX: Example: Multiple client sessions	408
AIX: Example: Simultaneous write operations	409
AIX: Example: Basic client store and server operations	410
AIX: Example: Conditions of extreme variation	411
AIX: Example: Full database backups	411
AIX: Example: Data deduplication	412
AIX: Active-log mirror space	416
AIX: Archive-failover log space	416
AIX: Monitoring space utilization for the database and recovery logs	416
AIX: Deleting installation rollback files	417
AIX: Deleting installation rollback files by using a graphical wizard	417
AIX: Deleting installation rollback files by using the command line	418
AIX: Server naming best practices	418
AIX: Installation directories for the IBM Spectrum Protect server	420
AIX: Installing the server components	420
AIX: Obtaining the installation package	420
AIX: Using the installation wizard	421
AIX: Using the console installation wizard	422
AIX: Using silent mode	423
AIX: Installing server language packages	424
AIX: Server language locales	424
AIX: Configuring a language package	425
AIX: Updating a language package	425
AIX: Taking the first steps after you install Version 8.1.5	425
AIX: Creating the user ID and directories for the server instance	426
AIX: Configuring the IBM Spectrum Protect server	427
AIX: Using the configuration wizard	428
AIX: Using the manual configuration steps	428
AIX: Creating the server instance	428
AIX: Configuring server and client communications on UNIX systems	430
AIX: Setting TCP/IP options	430
AIX: Setting shared memory options	431
AIX: Setting Secure Sockets Layer options	432
AIX: Formatting the database and log	432
AIX: Preparing the database manager for database backup	432
AIX: Configuring server options for server database maintenance	434
AIX: Starting the server instance	435
AIX: Verifying access rights and user limits	435
AIX: Starting the server from the instance user ID	436
AIX: Automatically starting servers	437
AIX: Starting the server in maintenance mode	438
AIX: Stopping the server	439
AIX: Registering licenses	439
AIX: Preparing the server for database backup operations	439
AIX: Running multiple server instances on a single system	440
AIX: Monitoring the server	440
AIX: Installing an IBM Spectrum Protect fix pack	441
AIX: Reverting from Version 8.1.5 to a previous server	443
AIX: Reference: DB2 commands for server databases	445
AIX: Uninstalling IBM Spectrum Protect	447
AIX: Uninstalling IBM Spectrum Protect by using a graphical wizard	448
AIX: Uninstalling IBM Spectrum Protect in console mode	448
AIX: Uninstalling IBM Spectrum Protect in silent mode	449
AIX: Uninstalling and reinstalling IBM Spectrum Protect	449

AIX: Uninstalling IBM Installation Manager	450
Linux: Installing the server	451
Linux: Planning to install the IBM Spectrum Protect server	451
Linux: What you should know first	452
Linux: Planning for optimal performance	452
Linux: Planning server hardware and operating system	452
Linux: Planning server database disks	455
Linux: Planning server recovery log disks	457
Linux: Planning container storage pools	458
Linux: Planning DISK or FILE storage pools	464
Linux: Planning storage technology	465
Linux: Installation best practices	467
Linux: Minimum system requirements for Linux systems	468
Linux: Minimum Linux x86_64 server requirements	469
Linux: Minimum Linux on System z server requirements	469
Linux: Minimum Linux on Power Systems (little endian) server requirements	469
Linux: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system	469
Linux: IBM Installation Manager	470
Linux: Worksheets for planning details for the server	471
Linux: Capacity planning	471
Linux: Database space requirements	472
Linux: Maximum number of files	472
Linux: Storage pool capacity	474
Linux: The database manager and temporary space	474
Linux: Recovery log space requirements	474
Linux: Active and archive log space	475
Linux: Example: Basic client-store operations	476
Linux: Example: Multiple client sessions	477
Linux: Example: Simultaneous write operations	478
Linux: Example: Basic client store and server operations	479
Linux: Example: Conditions of extreme variation	480
Linux: Example: Full database backups	480
Linux: Example: Data deduplication	481
Linux: Active-log mirror space	485
Linux: Archive-failover log space	485
Linux: Monitoring space utilization for the database and recovery logs	485
Linux: Deleting installation rollback files	486
Linux: Deleting installation rollback files by using a graphical wizard	486
Linux: Deleting installation rollback files by using the command line	487
Linux: Server naming best practices	487
Linux: Installation directories for the IBM Spectrum Protect server	489
Linux: Installing the server components	489
Linux: Obtaining the installation package	489
Linux: Using the installation wizard	490
Linux: Using the console installation wizard	491
Linux: Using silent mode	491
Linux: Installing server language packages	492
Linux: Server language locales	493
Linux: Configuring a language package	493
Linux: Updating a language package	494
Linux: Taking the first steps after you install Version 8.1.5	494
Linux: Tuning kernel parameters for Linux systems	495
Linux: Updating parameters	495
Linux: Suggested values	496
Linux: Creating the user ID and directories for the server instance	496
Linux: Configuring the IBM Spectrum Protect server	497

Linux: Using the configuration wizard	498
Linux: Using the manual configuration steps	498
Linux: Creating the server instance	498
Linux: Configuring server and client communications on UNIX systems	500
Linux: Setting TCP/IP options	500
Linux: Setting shared memory options	501
Linux: Setting Secure Sockets Layer options	502
Linux: Formatting the database and log	502
Linux: Preparing the database manager for database backup	502
Linux: Configuring server options for server database maintenance	504
Linux: Starting the server instance	505
Linux: Verifying access rights and user limits	505
Linux: Starting the server from the instance user ID	507
Linux: Automatically starting servers on Linux systems	507
Linux: Starting the server in maintenance mode	509
Linux: Stopping the server	509
Linux: Registering licenses	510
Linux: Preparing the server for database backup operations	510
Linux: Running multiple server instances on a single system	510
Linux: Monitoring the server	511
Linux: Installing an IBM Spectrum Protect fix pack	512
Linux: Reverting from Version 8.1.5 to a previous server	513
Linux: Reference: DB2 commands for server databases	515
Linux: Uninstalling IBM Spectrum Protect	518
Linux: Uninstalling IBM Spectrum Protect by using a graphical wizard	519
Linux: Uninstalling IBM Spectrum Protect in console mode	519
Linux: Uninstalling IBM Spectrum Protect in silent mode	519
Linux: Uninstalling and reinstalling IBM Spectrum Protect	520
Linux: Uninstalling IBM Installation Manager	521
Windows: Installing the server	521
Windows: Planning to install the IBM Spectrum Protect server	521
Windows: What you should know first	522
Windows: Planning for optimal performance	522
Windows: Planning server hardware and operating system	523
Windows: Planning server database disks	525
Windows: Planning server recovery log disks	527
Windows: Planning container storage pools	528
Windows: Planning DISK or FILE storage pools	534
Windows: Planning storage technology	535
Windows: Installation best practices	537
Windows: Minimum system requirements for Windows systems	538
Windows: IBM Installation Manager	539
Windows: Worksheets for planning details for the server	539
Windows: Capacity planning	540
Windows: Database space requirements	540
Windows: Maximum number of files	541
Windows: Storage pool capacity	542
Windows: The database manager and temporary space	543
Windows: Recovery log space requirements	543
Windows: Active and archive log space	543
Windows: Example: Basic client-store operations	544
Windows: Example: Multiple client sessions	545
Windows: Example: Simultaneous write operations	547
Windows: Example: Basic client store and server operations	548
Windows: Example: Conditions of extreme variation	548
Windows: Example: Full database backups	548

Windows: Example: Data deduplication	549
Windows: Active-log mirror space	553
Windows: Archive-failover log space	554
Windows: Monitoring space utilization for the database and recovery logs	554
Windows: Deleting installation rollback files	555
Windows: Deleting installation rollback files by using a graphical wizard	555
Windows: Deleting installation rollback files by using the command line	555
Windows: Server naming best practices	555
Windows: Installation directories for the IBM Spectrum Protect server	557
Windows: Installing the server components	557
Windows: Obtaining the installation package	557
Windows: Using the installation wizard	558
Windows: Using the console installation wizard	559
Windows: Using silent mode	559
Windows: Installing server language packages	560
Windows: Server language locales	561
Windows: Configuring a language package	561
Windows: Updating a language package	561
Windows: Taking the first steps after you install Version 8.1.5	562
Windows: Creating the user ID and directories for the server instance	563
Windows: Configuring the IBM Spectrum Protect server	564
Windows: Using the configuration wizard	564
Windows: Using the manual configuration steps	565
Windows: Creating the server instance	565
Windows: Configuring communications on Windows systems	566
Windows: Setting TCP/IP options	567
Windows: Setting Named Pipes options	567
Windows: Setting Secure Sockets Layer options	568
Windows: Formatting the database and log	568
Windows: Preparing the database manager for database backup	569
Windows: Configuring server options for server database maintenance	569
Windows: Starting the server instance on Windows systems	570
Windows: Configuring the server to start as a Windows service	571
Windows: Starting the server as a Windows service	572
Windows: Manually creating and configuring a Windows service	572
Windows: Starting the server in the foreground	573
Windows: Services associated with the server on Windows systems	573
Windows: Starting the server in maintenance mode	574
Windows: Stopping the server	575
Windows: Registering licenses	575
Windows: Preparing the server for database backup operations	575
Windows: Running multiple server instances on a single system	575
Windows: Monitoring the server	576
Windows: Installing an IBM Spectrum Protect fix pack	577
Windows: Reverting from Version 8.1.5 to a previous server	579
Windows: Reference: DB2 commands for server databases	581
Windows: Uninstalling IBM Spectrum Protect	584
Windows: Uninstalling IBM Spectrum Protect by using a graphical wizard	584
Windows: Uninstalling IBM Spectrum Protect in console mode	585
Windows: Uninstalling IBM Spectrum Protect in silent mode	585
Windows: Uninstalling and reinstalling IBM Spectrum Protect	586
Windows: Uninstalling IBM Installation Manager	587
Upgrading the server to V8.1	587
Upgrading to V8.1	588
Planning the upgrade	588
Preparing the system	589

Installing the server and verifying the upgrade	591
Upgrading the server in a clustered environment	595
Upgrading from V6.3 or V7.1 to V8.1.5 in a clustered environment for AIX with a shared database instance	596
Upgrading from V6.3 to V8.1.5 in a clustered environment for AIX with separate database instances	598
Upgrading to V8.1.5 in a clustered environment for Linux	600
Upgrading from V6.3 or V7.1 to V8.1.5 in a clustered environment for Windows	600
Installing and upgrading the Operations Center	602
Planning to install the Operations Center	603
System requirements for the Operations Center	603
Operations Center computer requirements	604
Hub and spoke server requirements	604
Tips for designing the hub and spoke server configuration	605
Tips for choosing a hub server	606
Operating system requirements	607
Web browser requirements	607
Language requirements	608
Requirements and limitations for IBM Spectrum Protect client management services	609
Administrator IDs that the Operations Center requires	610
IBM Installation Manager	611
Installation checklist	612
Installing the Operations Center	613
Obtaining the Operations Center installation package	614
Installing the Operations Center by using a graphical wizard	615
Installing the Operations Center in console mode	615
Installing the Operations Center in silent mode	616
Upgrading the Operations Center	617
Getting started with the Operations Center	617
Configuring the Operations Center	618
Designating the hub server	619
Adding a spoke server	619
Sending email alerts to administrators	620
Adding customized text to the login screen	621
Enabling REST services	622
Configuring for secure communication	622
Between the Operations Center and the hub server	623
Between the hub server and a spoke server	625
Resetting the password for the Operations Center truststore file	626
Starting and stopping the web server	627
Opening the Operations Center	628
Collecting diagnostic information with the client management service	628
Installing the client management service by using a graphical wizard	629
Installing the client management service in silent mode	630
Verifying the installation	631
Configuring the Operations Center to use the client management service	632
Starting and stopping the client management service	633
Uninstalling the client management service	633
Configuring the client management service for custom client installations	634
Troubleshooting the Operations Center installation	634
Graphical installation wizard cannot be started on an AIX system	634
Chinese, Japanese, or Korean fonts are displayed incorrectly	634
Uninstalling the Operations Center	635
Uninstalling the Operations Center by using a graphical wizard	635
Uninstalling the Operations Center in console mode	635
Uninstalling the Operations Center in silent mode	636
Rolling back to a previous version of the Operations Center	636
Configuring servers	637

Securing the server	639
Security concepts	640
Managing administrators	642
Changing password requirements	642
Securing IBM Spectrum Protect on the system	643
Restricting user access to the server	643
Limiting access through port restrictions	644
Protecting the storage environment against ransomware	644
Securing communications	645
SSL and TLS communication	646
Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL	648
Configuring the server to accept SSL connections	648
Configuring clients to communicate with the server by using SSL	649
Configuring the server to connect to another server by using SSL	650
Configuring the Operations Center to connect to the hub server by using SSL	651
Configuring the Data Protection for VMware vSphere GUI to communicate with the server by using SSL	651
Configuring a storage agent to use SSL	651
Configuring the client to connect to a storage agent by using SSL	652
Authenticating users by using an LDAP server	652
Replicating client data to another server	653
Replication compatibility	653
Enabling node replication	654
Protecting data in directory-container storage pools	655
Modifying replication settings	656
Setting different retention policies	656
Configuring clustered environments	657
Clustered environment overview	658
AIX clustered environment	658
Cluster requirements	659
PowerHA failover and failback	659
Installing and configuring PowerHA SystemMirror for AIX	660
Installing and configuring the cluster	660
Configuring on the primary node	660
Configuring on a secondary node with a shared DB2 instance	661
Configuring on a secondary node with a separate DB2 instance	662
Installing the server on a production node	663
Installing the client on a production node	663
Verifying the server configuration	664
Setting up the standby node	664
Defining the removable media storage devices	665
Configuring the cluster manager	665
Troubleshooting the PowerHA clustered environment	666
Linux clustered environment	666
Overview of a two-node clustered environment	667
Two-node shared disk topology	669
System Automation for Multiplatforms resource groups	670
Setting up a cluster	671
Prerequisites for configuring a cluster environment	671
Installing and configuring components	671
Installing server components	672
Configuring the primary node	672
Configuring the secondary node	673
Installing System Automation for Multiplatforms	674
Creating the label for the mount points	674
Installing and configuring System Automation for Multiplatforms	674
Preparing to activate the cluster nodes for the domain	675

Configuring volume group resources	675
Configuring resources that are not in a volume group	676
Activating the base policy	677
Adding mount points to directories	677
Configuring storage resources	678
Adding a storage pool	678
Deleting a storage pool	678
Deleting a mount point	679
Upgrading the server that is configured System Automation for Multiplatforms	679
Windows clustered environment	680
Microsoft Failover Cluster environment overview	680
Tape failover for nodes in a cluster	681
Planning for a clustered environment	682
Cluster configuration worksheet	682
Preparing Windows systems for a clustered environment	683
Configuring IBM Spectrum Protect in Microsoft Failover Cluster	683
Setting up IBM Spectrum Protect in a Microsoft Failover Cluster	684
Preparing a cluster resource group for a virtual server	684
Installing IBM Spectrum Protect in a Microsoft Failover Cluster	685
Initializing the server on the primary node	685
Verifying configuration in a Microsoft Failover Cluster	685
Testing failover	686
Maintaining the clustered environment	686
Migrating an existing server into a cluster	686
Adding a server by using backup and restore	687
Managing a virtual server on a cluster	687
Managing tape failover	687
Troubleshooting using the cluster log	688
Configuring clients	688
Adding clients	688
Selecting the client software and planning the installation	689
Specifying rules for backing up and archiving client data	690
Viewing policies	691
Editing policies	691
Scheduling backup and archive operations	692
Registering clients	693
Installing and configuring clients	694
Configuring the client to run scheduled operations	695
Configuring communications through a firewall	697
Scheduling client updates	697
Customizing policies	699
Policy concepts	699
Retention and expiration of backup versions	700
File expiration and expiration processing	701
Example: Retention when a policy uses only time controls	702
Example: Retention when a policy uses both version and time controls	703
Interactions among policy settings	704
Policy activation after updates	705
Customizing a policy	707
Creating a policy by copying an existing policy	708
Creating a policy domain	709
Controlling client operations through client option sets	709
Configuring storage	710
Types of storage pools	711
Data deduplication options	713
Configuring storage devices	714

Configuring a directory-container storage pool	714
Copying directory-container storage pools to tape	715
Rotating tape volumes offsite without DRM	717
Changing the volume reclamation threshold	717
Reclaiming tape volumes in container-copy storage pools	718
Determining whether to use container-copy storage pools for disaster protection	719
Configuring a cloud-container storage pool	721
Preparing for Amazon with S3 (off premises)	722
Preparing for an Amazon S3 compatible device	723
Preparing for Microsoft Azure (off premises)	724
Preparing for IBM Cloud Object Storage with Swift (off premises)	725
Preparing for IBM Cloud Object Storage with S3 (off premises)	725
Preparing for IBM Cloud Object Storage with S3 (on premises)	726
Preparing for OpenStack with Swift	728
Encrypting data for cloud-container storage pools	728
Defining a storage rule for cloud tiering	728
Reclaiming cloud containers	729
Optimizing performance for cloud object storage	730
Managing container storage pools	730
Converting a primary storage pool to a container storage pool	732
Cleaning up data in a source storage pool	733
Auditing a storage pool	734
Auditing a storage pool container	734
Storage system requirements and reducing the risk of data corruption	735
Monitoring storage solutions	736
Daily checklist	736
Periodic checklist	744
Verifying license compliance	749
Tracking system status by using email reports	750
Selecting, configuring, and using monitoring tools	751
Managing operations	753
Managing server operations	753
Stopping and starting the server	754
Stopping the server	754
Starting the server for maintenance or reconfiguration tasks	755
Managing inventory capacity	755
Managing memory and processor usage	757
Determining whether Aspera FASP can optimize data transfer in your environment	757
Planning to upgrade the server	759
Tuning scheduled activities	760
Managing client operations	760
Modifying the scope of a client backup	761
Evaluating errors in client error logs	761
Stopping and restarting the client acceptor	762
Resetting passwords	763
Decommissioning a client node	763
Deactivating data to free storage space	765
Managing client upgrades	766
Managing the Operations Center	767
Adding and removing spoke servers	767
Adding a spoke server	767
Removing a spoke server	768
Starting and stopping the web server	768
Restarting the initial configuration wizard	769
Changing the hub server	770
Restoring the configuration to the preconfiguration state	770

Configuring virtual tape libraries	771
Considerations for using virtual tape libraries	771
Storage capacity for virtual tape libraries	772
Drive configuration for virtual tape libraries	772
Adding a virtual tape library to your environment	773
Defining all drives and paths for a single library	773
Example: SCSI library or VTL with a single drive device type	774
Example: VTL or SCSI library with multiple drive device types	775
Protecting NAS file servers	777
NDMP requirements	778
Interfaces for NDMP operations	779
Data formats for NDMP backup operations	780
Storage pool types for NDMP operations	780
NDMP operations management	782
Managing NAS file server nodes	782
Managing data movers that are used in NDMP operations	783
Dedicating an IBM Spectrum Protect drive to NDMP operations	784
Storage pool management for NDMP operations	784
Managing tables of contents	785
Preventing inactive NDMP connections from closing	785
Enabling TCP keepalive	786
Specifying connection idle time (AIX, Linux, and Windows)	786
Configuring IBM Spectrum Protect for NDMP operations	786
In a nonclustered environment	786
Configuring an IBM Spectrum Protect policy for NDMP operations	788
Policies for backups initiated with an IBM Spectrum Protect server	789
Policies for backups initiated with the client interface	789
Determination of the NAS backup location	790
Tape libraries and drives for NDMP operations	791
Determining library drive usage when backing up to NAS-attached libraries	791
Configuring a tape library for NDMP operations	792
Attaching tape library robotics for NAS-attached libraries	794
Configuration 1: SCSI library connected to the IBM Spectrum Protect server	795
Configuration 2: SCSI library connected to the NAS file server	795
Configuration 3: 349x library connected to the IBM Spectrum Protect server	796
Configuration 4: ACSLS library connected to the IBM Spectrum Protect server	796
Registering NAS nodes with the IBM Spectrum Protect server	797
Defining a data mover for a NAS file server	797
Defining paths for NDMP operations	798
Defining paths to drives	798
Drives attached to a file server and the IBM Spectrum Protect server	798
Drives attached only to a file server	799
Obtaining names for devices attached to a file server	800
Defining paths to libraries	801
Scheduling NDMP operations	801
Defining virtual file spaces	802
Backing up data with the tape-to-tape function	802
Moving data with the tape-to-tape copy function	802
In a NetApp clustered environment	803
Configuring full cluster backups to tape devices	804
Configuring full cluster backups to an IBM Spectrum Protect server	806
Configuring partial cluster backups to an IBM Spectrum Protect server	807
Reconfiguring IBM Spectrum Protect to optimize clustered backups	808
Backing up and restoring NAS file servers using NDMP	810
NAS file servers: backups to a single IBM Spectrum Protect server	811
Backing up NDMP file servers to an IBM Spectrum Protect server	812

File-level backup and restore for NDMP operations _____	812
Interfaces for file-level restore operations _____	813
International characters for NetApp file servers _____	814
File-level restore operations from a directory-level backup image _____	814
Directory-level backup and restore operations _____	814
Directory-level backup and restore for NDMP operations _____	815
Backing up and restoring with snapshots _____	815
Backup and restore operations by using the NetApp SnapMirror to Tape feature _____	816
NDMP backup operations using Celerra file server-integrated checkpoints _____	816
Replicating NAS nodes _____	816
Data protection with the NetApp SnapLock feature _____	817
Reclamation and the SnapLock feature _____	818
Retention periods _____	818
Configuration of the SnapLock feature for event-based retention _____	820
Continuous data protection with the SnapLock feature _____	821
Setting up SnapLock volumes as IBM Spectrum Protect WORM FILE volumes _____	821
Repairing and recovering data _____	821
Repairing storage pools from a target replication server _____	822
Repairing storage pools from container-copy storage pool volumes _____	824
Repairing storage pools in an environment with both a replication server and container-copy storage pool volumes _____	826
Repairing storage pools on a target replication server _____	828
Repairing after a disaster _____	829
Repairing from container-copy storage pool volumes _____	829
Repairing from a target replication server _____	831
Repairing in an environment with both a replication server and container-copy storage pool volumes _____	832
Replacing a damaged container-copy storage pool tape volume _____	834
Server commands, options, and utilities _____	834
Managing the server from the command line _____	835
Issuing commands from the administrative client _____	836
Starting and stopping the administrative client _____	836
Monitoring server activities from the administrative client _____	837
Monitoring removable-media mounts from the administrative client _____	837
Processing individual commands from the administrative client _____	837
Processing a series of commands from the administrative client _____	838
Formatting output from commands _____	838
Saving command output to a specified location _____	838
Administrative client options _____	839
Issuing commands from the Operations Center _____	841
Issuing commands from the server console _____	841
Entering administrative commands _____	841
Reading syntax diagrams _____	842
Using continuation characters to enter long commands _____	845
Naming IBM Spectrum Protect objects _____	846
Using wildcard characters to specify object names _____	846
Specifying descriptions in keyword parameters _____	847
Controlling command processing _____	848
Server command processing _____	848
Stopping background processes _____	849
Performing tasks concurrently on multiple servers _____	849
Privilege classes for commands _____	851
Commands requiring system privilege _____	851
Commands requiring policy privilege _____	854
Commands requiring storage privilege _____	854
Commands requiring operator privilege _____	855
Commands any administrator can issue _____	856
Administrative commands _____	856

ACCEPT DATE (Accepts the current system date)	860
ACTIVATE POLICYSET (Activate a new policy set)	861
ASSIGN DEFMGMTCLASS (Assign a default management class)	862
AUDIT commands	863
AUDIT CONTAINER commands	863
Cloud-container audit	863
Directory-container audit	868
AUDIT LDAPDIRECTORY (Audit an LDAP directory server)	871
AUDIT LIBRARY (Audit volume inventories in an automated library)	873
AUDIT LIBVOLUME (Verify database information for a tape volume)	875
AUDIT LICENSES (Audit server storage usage)	876
AUDIT VOLUME (Verify database information for a storage pool volume)	877
BACKUP commands	881
BACKUP DB (Back up the database)	881
BACKUP DEVCONFIG (Create backup copies of device configuration information)	885
BACKUP NODE (Back up a NAS node)	887
BACKUP STGPOOL (Back up primary storage pool data to a copy storage pool)	890
BACKUP VOLHISTORY (Save sequential volume history information)	893
BEGIN EVENTLOGGING (Begin logging events)	894
CANCEL commands	895
CANCEL EXPIRATION (Cancel an expiration process)	896
CANCEL EXPORT (Delete a suspended export operation)	896
CANCEL PROCESS (Cancel an administrative process)	897
CANCEL REPLICATION (Cancel node replication processes)	899
CANCEL REQUEST (Cancel one or more mount requests)	899
CANCEL RESTORE (Cancel a restartable restore session)	900
CANCEL SESSION (Cancel one or more client sessions)	901
CHECKIN LIBVOLUME (Check a storage volume into a library)	902
CHECKOUT LIBVOLUME (Check a storage volume out of a library)	907
CLEAN DRIVE (Clean a drive)	911
COMMIT (Control committing of commands in a macro)	912
CONVERT STGPOOL (Convert a storage pool to a container storage pool)	913
COPY commands	914
COPY ACTIVATEDATA (Copy active backup data from a primary storage pool to an active-data pool)	915
COPY CLOPTSET (Copy a client option set)	917
COPY DOMAIN (Copy a policy domain)	918
COPY MGMTCLASS (Copy a management class)	919
COPY POLICYSET (Copy a policy set)	920
COPY PROFILE (Copy a profile)	921
COPY SCHEDULE (Copy a client or an administrative command schedule)	922
COPY SCHEDULE (Create a copy of a schedule for client operations)	922
COPY SCHEDULE (Create a copy of a schedule for administrative operations)	923
COPY SCRIPT (Copy an IBM Spectrum Protect script)	924
COPY SERVERGROUP (Copy a server group)	925
DEACTIVATE DATA (Deactivate data for a client node)	926
DECOMMISSION commands	927
DECOMMISSION NODE (Decommission an application or system)	928
DECOMMISSION VM (Decommission a virtual machine)	929
DEFINE commands	930
DEFINE ALERTTRIGGER (Define an alert trigger)	931
DEFINE ASSOCIATION (Associate client nodes with a schedule)	933
DEFINE BACKUPSET (Define a backup set)	934
DEFINE CLIENTACTION (Define a one-time client action)	937
DEFINE CLIENTOPT (Define an option to an option set)	941
DEFINE CLOPTSET (Define a client option set name)	943
DEFINE COLLOGGROUP (Define a collocation group)	944

DEFINE COLLOCMEMBER	945
DEFINE COPYGROUP (Define a copy group)	947
DEFINE COPYGROUP (Define a backup copy group)	948
DEFINE COPYGROUP (Define an archive copy group)	951
DEFINE DATAMOVER (Define a data mover)	954
DEFINE DEVCLASS (Define a device class)	956
3590	956
3592	959
4MM	965
8MM	968
Centera	972
DLT	974
Ecartridge	978
File	983
Generictape	985
LTO	987
NAS	992
Removablefile	993
Server	995
VolSafe	997
DEFINE DEVCLASS - z/OS media server (Define device class for z/OS media server)	1000
3590, for z/OS media server	1000
3592, for z/OS media server	1004
ECARTRIDGE, for z/OS media server	1008
FILE, for z/OS media server	1012
DEFINE DOMAIN (Define a new policy domain)	1014
DEFINE DRIVE (Define a drive to a library)	1016
DEFINE EVENTSERVER (Define a server as the event server)	1019
DEFINE GRPMEMBER (Add a server to a server group)	1020
DEFINE LIBRARY (Define a library)	1021
349X	1022
ACSLs	1025
EXTERNAL	1027
FILE	1028
MANUAL	1029
SCSI	1030
SHARED	1033
VTL	1033
ZOSMEDIA	1036
DEFINE MACHINE (Define machine information for disaster recovery)	1037
DEFINE MACHNODEASSOCIATION (Associate a node with a machine)	1038
DEFINE MGMTCLASS (Define a management class)	1039
DEFINE NODEGROUP (Define a node group)	1041
DEFINE NODEGROUPMEMBER (Define node group member)	1042
DEFINE PATH (Define a path)	1043
Destination is a drive	1043
Destination is a library	1048
Destination is a ZOSMEDIA library	1051
DEFINE POLICYSET (Define a policy set)	1051
DEFINE PROFASSOCIATION (Define a profile association)	1052
DEFINE PROFILE (Define a profile)	1056
DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)	1057
DEFINE RECOVERYMEDIA (Define recovery media)	1058
DEFINE SCHEDULE (Define a client or an administrative command schedule)	1059
DEFINE SCHEDULE (Define a client schedule)	1060
DEFINE SCHEDULE (Define a schedule for an administrative command)	1070

DEFINE SCRATCHPADENTRY (Define a scratch pad entry)	1077
DEFINE SCRIPT (Define an IBM Spectrum Protect script)	1078
DEFINE SERVER (Define a server for server-to-server communications)	1080
DEFINE SERVERGROUP (Define a server group)	1086
DEFINE SPACETRIGGER (Define the space trigger)	1087
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	1089
DEFINE STGPOOL (Define a storage pool)	1092
Cloud-container storage pool	1093
Directory-container storage pool	1098
Container-copy storage pool	1101
Primary random-access pool	1104
Primary sequential-access pool	1111
Copy pool	1124
Active-data pool	1130
DEFINE STGPOOLDIRECTORY (Define a storage pool directory)	1136
DEFINE STGRULE (Define a storage rule)	1137
DEFINE STGRULE (Define a rule for auditing storage pools)	1137
DEFINE STGRULE (Define a rule for generating data deduplication statistics)	1139
DEFINE STGRULE (Define a rule for reclaiming cloud containers)	1142
DEFINE STGRULE (Define a storage rule for tiering)	1143
DEFINE SUBSCRIPTION (Define a profile subscription)	1145
DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)	1146
DEFINE VOLUME (Define a volume in a storage pool)	1148
DELETE commands	1154
DELETE ALERTTRIGGER (Remove a message from an alert trigger)	1154
DELETE ASSOCIATION (Delete the node association to a schedule)	1155
DELETE BACKUPSET (Delete a backup set)	1156
DELETE CLIENTOPT (Delete an option in an option set)	1160
DELETE CLOPTSET (Delete a client option set)	1161
DELETE COLLOCGROUP (Delete a collocation group)	1161
DELETE COLLOCMEMBER (Delete collocation group member)	1162
DELETE COPYGROUP (Delete a backup or archive copy group)	1165
DELETE DATAMOVER (Delete a data mover)	1166
DELETE DEDUPSTATS (Delete data deduplication statistics)	1166
DELETE DEVCLASS (Delete a device class)	1169
DELETE DOMAIN (Delete a policy domain)	1170
DELETE DRIVE (Delete a drive from a library)	1171
DELETE EVENT (Delete event records)	1171
DELETE EVENTSERVER (Delete the definition of the event server)	1173
DELETE FILESPACE (Delete client node data from the server)	1173
DELETE GRPMEMBER (Delete a server from a server group)	1176
DELETE LIBRARY (Delete a library)	1177
DELETE MACHINE (Delete machine information)	1178
DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)	1179
DELETE MGMTCLASS (Delete a management class)	1180
DELETE NODEGROUP (Delete a node group)	1180
DELETE NODEGROUPMEMBER (Delete node group member)	1181
DELETE PATH (Delete a path)	1182
DELETE POLICYSET (Delete a policy set)	1183
DELETE PROFASSOCIATION (Delete a profile association)	1184
DELETE PROFILE (Delete a profile)	1186
DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)	1188
DELETE RECOVERYMEDIA (Delete recovery media)	1188
DELETE SCHEDULE (Delete a client or an administrative command schedule)	1189
DELETE SCHEDULE (Delete a client schedule)	1189
DELETE SCHEDULE (Delete an administrative schedule)	1190

DELETE SCRATCHPADENTRY (Delete a scratch pad entry)	1190
DELETE SCRIPT (Delete command lines from a script or delete the entire script)	1191
DELETE SERVER (Delete a server definition)	1192
DELETE SERVERGROUP (Delete a server group)	1193
DELETE SPACETRIGGER (Delete the storage pool space triggers)	1193
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	1194
DELETE STGPOOL (Delete a storage pool)	1195
DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)	1196
DELETE STGRULE (Delete storage rules for storage pools)	1197
DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)	1198
DELETE SUBSCRIPTION (Delete a profile subscription)	1199
DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)	1199
DELETE VOLHISTORY (Delete sequential volume history information)	1200
DELETE VOLUME (Delete a storage pool volume)	1204
DISABLE commands	1206
DISABLE EVENTS (Disable events for event logging)	1206
DISABLE REPLICATION (Prevent outbound replication processing on a server)	1209
DISABLE SESSIONS (Prevent new sessions from accessing IBM Spectrum Protect)	1209
DISMOUNT command	1211
DISPLAY OBJNAME (Display a full object name)	1211
ENABLE commands	1212
ENABLE EVENTS (Enable server or client events for logging)	1212
ENABLE REPLICATION (Allow outbound replication processing on a server)	1214
ENABLE SESSIONS (Resume user activity on the server)	1215
ENCRYPT STGPOOL (Encrypt data in a storage pool)	1217
END EVENTLOGGING (Stop logging events)	1218
EXPIRE INVENTORY (Manually start inventory expiration processing)	1219
EXPORT commands	1222
EXPORT ADMIN (Export administrator information)	1222
EXPORT ADMIN (Export administrator definitions to sequential media)	1224
EXPORT ADMIN (Export administrator information directly to another server)	1226
EXPORT NODE (Export client node information)	1228
EXPORT NODE (Export node definitions to sequential media)	1230
EXPORT NODE (Export node definitions or file data directly to another server)	1236
EXPORT POLICY (Export policy information)	1243
EXPORT POLICY (Export policy information to sequential media)	1244
EXPORT POLICY (Export a policy directly to another server)	1246
EXPORT SERVER (Export server information)	1248
EXPORT SERVER (Export a server to sequential media)	1249
EXPORT SERVER (Export server control information and client file data to another server)	1255
EXTEND DBSPACE (Increase space for the database)	1261
GENERATE commands	1263
GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)	1263
GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)	1269
GENERATE DEDUPSTATS (Generate data deduplication statistics)	1270
GRANT commands	1273
GRANT AUTHORITY (Add administrator authority)	1273
GRANT PROXYNODE (Grant proxy authority to a client node)	1276
HALT (Shut down the server)	1276
HELP (Get help on commands and error messages)	1277
IDENTIFY DUPLICATES (Identify duplicate data in a storage pool)	1279
IMPORT commands	1282
IMPORT ADMIN (Import administrator information)	1282
IMPORT NODE (Import client node information)	1284
IMPORT POLICY (Import policy information)	1290
IMPORT SERVER (Import server information)	1292

INSERT MACHINE (Insert machine characteristics information or recovery instructions)	1297
ISSUE MESSAGE (Issue a message from a server script)	1298
LABEL LIBVOLUME (Label a library volume)	1299
LOAD DEFALERTTRIGGERS (Load the default set of alert triggers)	1304
LOCK commands	1305
LOCK ADMIN (Lock out an administrator)	1305
LOCK NODE (Lock out a client node)	1306
LOCK PROFILE (Lock a profile)	1307
MACRO (Invoke a macro)	1308
MIGRATE STGPOOL (Migrate storage pool to next storage pool)	1309
MOVE commands	1311
MOVE CONTAINER (Move a container)	1311
MOVE DATA (Move files on a storage pool volume)	1313
MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)	1316
MOVE GRPMEMBER (Move a server group member)	1329
MOVE MEDIA (Move sequential-access storage pool media)	1329
MOVE NODEDATA (Move data by node in a sequential access storage pool)	1335
File spaces for one or more nodes or a collocation group	1336
Selected file spaces of a single node	1338
NOTIFY SUBSCRIBERS (Notify managed servers to update profiles)	1341
PERFORM LIBACTION (Define or delete all drives and paths for a library)	1342
PING SERVER (Test the connection between servers)	1346
PREPARE (Create a recovery plan file)	1346
PROTECT STGPOOL (Protect data that belongs to a storage pool)	1352
QUERY commands	1357
QUERY ACTLOG (Query the activity log)	1358
QUERY ADMIN (Display administrator information)	1363
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	1367
QUERY ALERTSTATUS (Query the status of an alert)	1368
QUERY ASSOCIATION (Query client node associations with a schedule)	1372
QUERY AUDITOCAPACITY (Query client node storage utilization)	1373
QUERY BACKUPSET (Query a backup set)	1374
QUERY BACKUPSETCONTENTS (Query contents of a backup set)	1378
QUERY CLEANUP (Query the cleanup that is required in a source storage pool)	1380
QUERY CLOPTSET (Query a client option set)	1381
QUERY COLLOGGROUP (Query a collocation group)	1383
QUERY CONTAINER (Display container information)	1385
QUERY CONTENT (Query the contents of a storage pool volume)	1388
QUERY CONVERSION (Query conversion status of a storage pool)	1394
QUERY COPYGROUP (Query copy groups)	1395
QUERY DAMAGED (Query damaged in a directory-container or cloud-container storage pool)	1398
QUERY DATAMOVER (Display data mover definitions)	1401
QUERY DB (Display database information)	1404
QUERY DBSPACE (Display database storage space)	1406
QUERY DEDUPSTATS (Query data deduplication statistics)	1407
QUERY DEVCLASS (Display information on one or more device classes)	1413
QUERY DIRSPACE (Query storage utilization of FILE directories)	1417
QUERY DOMAIN (Query a policy domain)	1418
QUERY DRIVE (Query information about a drive)	1420
QUERY DRMEDIA (Query disaster recovery media)	1423
QUERY DRMSTATUS (Query disaster recovery manager system parameters)	1430
QUERY ENABLED (Query enabled events)	1432
QUERY EVENT (Query scheduled and completed events)	1434
QUERY EVENT (Display client schedules)	1434
QUERY EVENT (Display administrative event schedules)	1440
QUERY EVENTRULES (Query rules for server or client events)	1443

QUERY EVENTSERVER (Query the event server)	1445
QUERY EXPORT (Query for active or suspended export operations)	1445
QUERY EXTENTUPDATES (Query updated data extents)	1450
QUERY FILESPACE (Query one or more file spaces)	1451
QUERY FSCOUNTS (Query number of objects)	1456
QUERY LIBRARY (Query a library)	1458
QUERY LIBVOLUME (Query a library volume)	1460
QUERY LICENSE (Display license information)	1462
QUERY LOG (Display information about the recovery log)	1465
QUERY MACHINE (Query machine information)	1467
QUERY MEDIA (Query sequential-access storage pool media)	1469
QUERY MGMTCLASS (Query a management class)	1474
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	1476
QUERY MONITORSTATUS (Query the monitoring status)	1478
QUERY MOUNT (Display information on mounted sequential access volumes)	1482
QUERY NASBACKUP (Query NAS backup images)	1483
QUERY NODE (Query nodes)	1487
QUERY NODEDATA (Query client data in volumes)	1496
QUERY NODEGROUP (Query a node group)	1498
QUERY OCCUPANCY (Query client file spaces in storage pools)	1499
QUERY OPTION (Query server options)	1502
QUERY PATH (Display a path definition)	1504
QUERY POLICYSET (Query a policy set)	1507
QUERY PROCESS (Query one or more server processes)	1509
QUERY PROFILE (Query a profile)	1513
QUERY PROTECTSTATUS (Query the status of storage pool protection)	1515
QUERY PROXYNODE (Query proxy authority for a client node)	1517
QUERY PVUESTIMATE (Display processor value unit estimate)	1518
QUERY RECOVERYMEDIA (Query recovery media)	1521
QUERY REPLICATION (Query node replication processes)	1523
QUERY REPLNODE (Display information about replication status for a client node)	1531
QUERY REPLRULE (Query replication rules)	1534
QUERY REPLSERVER (Query a replication server)	1535
QUERY REQUEST (Query one or more pending mount requests)	1537
QUERY RESTORE (Query restartable restore sessions)	1538
QUERY RPFCONTENT (Query recovery plan file contents stored on a target server)	1540
QUERY RPFFILE (Query recovery plan file information stored on a target server)	1541
QUERY SAN (Query the devices on the SAN)	1543
QUERY SCHEDULE (Query schedules)	1545
QUERY SCHEDULE (Query client schedules)	1546
QUERY SCHEDULE (Query an administrative schedule)	1549
QUERY SCRATCHPADENTRY (Query a scratch pad entry)	1551
QUERY SCRIPT (Query IBM Spectrum Protect scripts)	1552
QUERY SERVER (Query a server)	1555
QUERY SERVERGROUP (Query a server group)	1558
QUERY SESSION (Query client sessions)	1559
QUERY SHREDSTATUS (Query shredding status)	1563
QUERY SPACETRIGGER (Query the space triggers)	1564
QUERY STATUS (Query system parameters)	1565
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	1574
QUERY STGPOOL (Query storage pools)	1576
QUERY STGPOOLDIRECTORY (Query a storage pool directory)	1590
QUERY STGRULE (Display storage rule information)	1592
QUERY SUBSCRIBER (Display subscriber information)	1596
QUERY SUBSCRIPTION (Display subscription information)	1597
QUERY SYSTEM (Query the system configuration and capacity)	1598

QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)	1600
QUERY TOC (Display table of contents for a backup image)	1600
QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)	1602
QUERY VOLHISTORY (Display sequential volume history information)	1603
QUERY VOLUME (Query storage pool volumes)	1609
QUIT (End the interactive mode of the administrative client)	1616
RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)	1616
RECONCILE VOLUMES (Reconcile differences in the virtual volume definitions)	1618
REGISTER commands	1620
REGISTER ADMIN (Register an administrator ID)	1620
REGISTER LICENSE (Register a new license)	1624
REGISTER NODE (Register a node)	1625
REMOVE commands	1639
REMOVE ADMIN (Delete an administrative user ID)	1639
REMOVE DAMAGED (Remove damaged data from a source storage pool)	1640
REMOVE NODE (Delete a node or an associated machine node)	1641
REMOVE REPLNODE (Remove a client node from replication)	1642
REMOVE REPLSERVER (Remove a replication server)	1643
RENAME commands	1644
RENAME ADMIN (Rename an administrator)	1644
RENAME FILESPACE (Rename a client file space on the server)	1645
RENAME NODE (Rename a node)	1648
RENAME SCRIPT (Rename an IBM Spectrum Protect script)	1649
RENAME SERVERGROUP (Rename a server group)	1650
RENAME STGPOOL (Change the name of a storage pool)	1650
REPAIR STGPOOL (Repair a directory-container storage pool)	1651
REPLICATE NODE (Replicate data in file spaces that belong to a client node)	1653
REPLY (Allow a request to continue processing)	1661
RESET PASSEXP (Reset password expiration)	1662
RESTART EXPORT (Restart a suspended export operation)	1663
RESTORE commands	1664
RESTORE NODE (Restore a NAS node)	1664
RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)	1668
RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)	1671
REVOKE commands	1674
REVOKE AUTHORITY (Remove administrator authority)	1674
REVOKE PROXYNODE (Revoke proxy authority for a client node)	1676
ROLLBACK (Rollback uncommitted changes in a macro)	1677
RUN (Run an IBM Spectrum Protect script)	1678
SELECT (Perform an SQL query of the IBM Spectrum Protect database)	1680
SET commands	1688
SET ACCOUNTING (Set accounting records on or off)	1689
SET ACTLOGRETENTION (Set the retention period or the size of the activity log)	1690
SET ALERTACTIVEDURATION (Set the duration of an active alert)	1691
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	1692
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	1693
SET ALERTEMAILFROMADDR (Set the email address of the sender)	1693
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	1694
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	1695
SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	1695
SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	1696
SET ALERTMONITOR (Set the alert monitor to on or off)	1697
SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	1698
SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)	1698
SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)	1699
SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)	1701

SET CLIENTACTDURATION (Set the duration period for the client action)	1702
SET CONFIGMANAGER (Specify a configuration manager)	1703
SET CONFIGREFRESH (Set managed server configuration refresh)	1704
SET CONTEXTMESSAGING (Set message context reporting on or off)	1705
SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)	1705
SET CROSSDEFINE (Specifies whether to cross-define servers)	1706
SET DBRECOVERY (Set the device class for automatic backups)	1706
SET DEDUPVERIFICATIONLEVEL (Set the percentage of extents to verify)	1708
SET DEFAULTAUTHENTICATION (Set the default authentication method for REGISTER NODE and REGISTER ADMIN commands)	1710
SET DEPLOYPKGMR (Enable the deployment package manager)	1710
SET DEPLOYREPOSITORY (Set the download path for client deployment packages)	1711
SET DEPLOYMAXPKGS (Set the maximum number of client deployment packages to store)	1712
SET DISSIMILARPOLICIES (Enable the policies on the target replication server to manage replicated data)	1713
SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)	1714
SET DRMCHECKLABEL (Specify label checking)	1714
SET DRMCMDFILENAME (Specify the name of a file to contain commands)	1715
SET DRMCOPYCONTAINERSTGPOOL (Specify the container-copy storage pools to be processed by DRM commands)	1716
SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)	1717
SET DRMCOURIERNAME (Specify the courier name)	1718
SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)	1718
SET DRMFILEPROCESS (Specify file processing)	1719
SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)	1720
SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)	1722
SET DRMPPLANPREFIX (Specify a prefix for recovery plan file names)	1722
SET DRMPPLANPOSTFIX (Specify replacement volume names)	1724
SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)	1725
SET DRMRPFEXPIREDAYS (Set criteria for recovery plan file expiration)	1726
SET DRMVaultNAME (Specify the vault name)	1727
SET EVENTRETENTION (Set the retention period for event records)	1727
SET FAILOVERHLADDRESS (Set a failover high level address)	1728
SET INVALIDPWLIMIT (Set the number of invalid logon attempts)	1729
SET LDAPPASSWORD (Set the LDAP password for the server)	1730
SET LDAPUSER (Specify an ID for an LDAP directory server)	1731
SET LICENSEAUDITPERIOD (Set license audit period)	1731
SET MAXCMDRETRIES (Set the maximum number of command retries)	1732
SET MAXSCHEDULESESSIONS (Set maximum scheduled sessions)	1733
SET MINPWLENGTH (Set minimum password length)	1734
SET MONITOREDSEVERGROUP (Set the group of monitored servers)	1735
SET MONITORINGADMIN (Set the name of the monitoring administrator)	1735
SET NODEATRISKINTERVAL (Specifies at-risk mode for an individual node)	1736
SET PASSEXP (Set password expiration date)	1738
SET PRODUCTOFFERING (Set the product offering that is licensed to your enterprise)	1739
SET QUERYSCHEDPERIOD (Set query period for polling client nodes)	1740
SET RANDOMIZE (Set randomization of scheduled start times)	1741
SET REPLRECOVERDAMAGED (Specify whether damaged files are recovered from a replication server)	1742
SET REPLRETENTION (Set the retention period for replication records)	1744
SET REPLSERVER (Set the target replication server)	1745
SET RETRYPERIOD (Set time between retry attempts)	1746
SET SCHEDMODES (Select a central scheduling mode)	1746
SET SCRATCHPADRETENTION (Set scratch pad retention time)	1747
SET SERVERHLADDRESS (Set the high-level address of a server)	1748
SET SERVERLLADDRESS (Set the low-level address of a server)	1749
SET SERVERNAME (Specify the server name)	1749
SET SERVERPASSWORD (Set password for server)	1750
SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)	1751

SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	1752
SET STATUSMONITOR (Specifies whether to enable status monitoring)	1753
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	1754
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	1755
SET SUBFILE (Set subfile backup for client nodes)	1756
SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)	1757
SET TAPEALERTMSG (Set tape alert messages on or off)	1758
SET TOCLOADRETENTION (Set load retention period for table of contents)	1759
SET VMATRISKINTERVAL (Specifies the at-risk mode for an individual VM filespace)	1759
SETOPT (Set a server option for dynamic update)	1761
SHRED DATA (Shred data)	1762
SUSPEND EXPORT (Suspend a currently running export operation)	1764
UNLOCK commands	1765
UNLOCK ADMIN (Unlock an administrator)	1765
UNLOCK NODE (Unlock a client node)	1766
UNLOCK PROFILE (Unlock a profile)	1767
UPDATE commands	1767
UPDATE ALERTTRIGGER (Update a defined alert trigger)	1768
UPDATE ALERTSTATUS (Update the status of an alert)	1770
UPDATE ADMIN (Update an administrator)	1771
UPDATE BACKUPSET (Update a retention value assigned to a backup set)	1775
UPDATE CLIENTOPT (Update a client option sequence number)	1778
UPDATE CLOPTSET (Update a client option set description)	1779
UPDATE COLLOCGROUP (Update a collocation group)	1780
UPDATE COPYGROUP (Update a copy group)	1781
UPDATE COPYGROUP (Update a backup copy group)	1781
UPDATE COPYGROUP (Update a defined archive copy group)	1784
UPDATE DATAMOVER (Update a data mover)	1786
UPDATE DEVCLASS (Update the attributes of a device class)	1787
3590	1788
3592	1791
4MM	1796
8MM	1799
Centera	1803
DLT	1805
Ecartridge	1809
File	1813
Generictape	1817
LTO	1818
NAS	1823
Removablefile	1825
Server	1826
VolSafe	1828
UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server)	1830
3590, for z/OS media server	1831
3592, for z/OS media server	1834
ECARTRIDGE, for z/OS media server	1838
FILE, for z/OS media server	1841
UPDATE DOMAIN (Update a policy domain)	1843
UPDATE DRIVE (Update a drive)	1845
UPDATE FILESPACE (Update file-space node-replication rules)	1848
UPDATE LIBRARY (Update a library)	1851
349X	1852
ACSLs	1854
EXTERNAL	1856
FILE	1856

MANUAL	1857
SCSI	1858
SHARED	1860
VTL	1861
UPDATE LIBVOLUME (Change the status of a storage volume)	1863
UPDATE MACHINE (Update machine information)	1864
UPDATE MGMTCLASS (Update a management class)	1865
UPDATE NODE (Update node attributes)	1867
UPDATE NODEGROUP (Update a node group)	1881
UPDATE PATH (Change a path)	1882
Destination is a drive	1882
Destination is a library	1886
Destination is a ZOSMEDIA library	1888
UPDATE POLICYSET (Update a policy set description)	1889
UPDATE PROFILE (Update a profile description)	1890
UPDATE RECOVERYMEDIA (Update recovery media)	1891
UPDATE REPLRULE (Update replication rules)	1892
UPDATE SCHEDULE (Update a schedule)	1893
UPDATE SCHEDULE (Update a client schedule)	1894
UPDATE SCHEDULE (Update an administrative schedule)	1903
UPDATE SCRATCHPADENTRY (Update a scratch pad entry)	1910
UPDATE SCRIPT (Update an IBM Spectrum Protect script)	1911
UPDATE SERVER (Update a server defined for server-to-server communications)	1913
UPDATE SERVERGROUP (Update a server group description)	1917
UPDATE SPACETRIGGER (Update the space triggers)	1918
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	1920
UPDATE STGPOOL (Update a storage pool)	1922
Cloud-container storage pool	1924
Directory-container storage pool	1927
Container-copy storage pool	1930
Primary random-access pool	1932
Primary sequential-access pool	1939
Copy pool	1950
Active-data pool	1955
UPDATE STGPOOLDIRECTORY (Update a storage pool directory)	1959
UPDATE STGRULE (Update a storage rule)	1961
UPDATE STGRULE (Update a rule for auditing a storage pool)	1962
UPDATE STGRULE (Update a storage rule for generating data deduplication statistics)	1963
UPDATE STGRULE (Update a storage rule for reclaiming cloud containers)	1966
UPDATE STGRULE (Update a storage rule for tiering)	1967
UPDATE VIRTUALFSMAPPING (Update a virtual file space mapping)	1968
UPDATE VOLHISTORY (Update sequential volume history information)	1970
UPDATE VOLUME (Change a storage pool volume)	1971
VALIDATE commands	1974
VALIDATE ASPERA (Validate an Aspera FASP configuration)	1974
VALIDATE CLOUD (Validate cloud credentials)	1977
VALIDATE LANFREE (Validate LAN-Free paths)	1979
VALIDATE POLICYSET (Verify a policy set)	1980
VALIDATE REPLICATION (Validate replication for a client node)	1982
VALIDATE REPLPOLICY (Verify the policies on the target replication server)	1985
VARY (Bring a random access volume online or offline)	1987
Server options	1988
Modifying server options	1995
Types of server options	1995
Server communication options	1996
Server storage options	1997

Client-server options	1998
Date, number, time, and language options	1999
Database options	1999
Data transfer options	2000
Message options	2000
Event logging options	2000
Security options and licensing options	2001
Miscellaneous options	2001
3494SHARED	2002
ACSACCESSID	2002
ACSLOCKDRIVE	2003
ACSQUICKINIT	2003
ACSTIMEOUTX	2004
ACTIVELOGDIRECTORY	2004
ACTIVELOGSIZE	2005
ADMINCOMMTIMEOUT	2005
ADMINIDLETIMEOUT	2006
ADMINONCLIENTPORT	2006
ADSMGROUPNAME	2006
ALIASHALT	2007
ALLOWDESAUTH	2007
ALLOWREORGINDEX	2008
ALLOWREORGTABLE	2008
ARCHFAILOVERLOGDIRECTORY	2009
ARCHLOGCOMPRESS	2009
ARCHLOGDIRECTORY	2010
ARCHLOGUSEDTHRESHOLD	2010
ASSISTVCRRECOVERY	2010
AUDITSTORAGE	2011
BACKUPINITIATIONROOT	2011
CHECKTAPEPOS	2012
CLIENTDEDUPTXNLIMIT	2013
CLIENTDEPLOYCATALOGURL	2014
CLIENTDEPLOYUSELOCALCATALOG	2014
COMMMETHOD	2015
COMMTIMEOUT	2015
CONTAINERRESOURCESTIMEOUT	2016
DATEFORMAT	2016
DBDIAGLOGSIZE	2017
DBDIAGPATHFSTHRESHOLD	2018
DBMEMPERCENT	2018
DBMTCPPORT	2019
DEDUPREQUIRESBACKUP	2019
DEDUPTIER2FILESIZE	2020
DEDUPTIER3FILESIZE	2020
DEVCONFIG	2021
DISABLEREORGTABLE	2022
DISABLESCHEDS	2022
DISPLAYLFINFO	2022
DNSLOOKUP	2023
DRIVEACQUIRERETRY	2024
ENABLENASDEDUP	2024
EVENTSERVER	2025
EXPINTERVAL	2025
EXPQUIET	2026
FASPBEGPORT	2026

FASPENPORT	2026
FASPTARGETRATE	2027
FFDCLOGLEVEL	2028
FFDCLOGNAME	2028
FFDCMAXLOGSIZE	2029
FFDCNUMLOGS	2029
FILEEXIT	2030
FILETEXTEXIT	2030
FIPSMODE	2031
FSUSEDTHRESHOLD	2031
IDLETIMEOUT	2032
KEEPALIVE	2032
KEEPALIVETIME	2033
KEEPALIVEINTERVAL	2033
LANGUAGE	2034
LDAPCACHEDURATION	2036
LDAPURL	2037
MAXSESSIONS	2038
MESSAGEFORMAT	2038
MIRRORLOGDIRECTORY	2038
MOVEBATCHSIZE	2039
MOVESIZETHRESH	2039
MSGINTERVAL	2040
NAMEDPIPENAME	2040
NDMPCONNECTIONTIMEOUT	2040
NDMPCONTROLPORT	2041
NDMPENABLEKEEPALIVE	2041
NDMPKEEPIDLEMINUTES	2042
NDMPPORTRANGE	2042
NDMPREFDATAINTERFACE	2043
NOPREEMPT	2043
NORETRIEVEDATE	2044
NPAUDITFAILURE	2044
NPAUDITSUCCESS	2045
NPBUFFERSIZE	2045
NUMBERFORMAT	2045
NUMOPENVOLSALLOWED	2046
PUSHSTATUS	2047
QUERYAUTH	2047
RECLAIMDELAY	2048
RECLAIMPERIOD	2048
REORGBEGINTIME	2049
REORGDURATION	2049
REPORTRETRIEVE	2050
REPLBATCHSIZE	2050
REPLSIZETHRESH	2051
REQSYSAUTHOUTFILE	2051
RESOURCE TIMEOUT	2052
RESTHTTPSPORT	2052
RESTOREINTERVAL	2053
RETENTIONEXTENSION	2053
SANDISCOVERY	2054
SANDISCOVERYTIMEOUT	2055
SANREFRESHTIME	2055
SEARCHMPQUEUE	2056
SECUREPIPES	2056

SERVERDEDUPTXNLIMIT	2056
SHMPORT	2057
SHREDDING	2058
SNMPHEARTBEATINTERVAL	2058
SNMPMESSAGECATEGORY	2058
SNMPSUBAGENT	2059
SNMPSUBAGENTHOST	2060
SNMPSUBAGENTPORT	2060
SSLFIPSMODE	2060
SSLINITTIMEOUT	2061
SSLTCPADMINPORT	2061
SSLTCPPOINT	2062
TCPADMINPORT	2063
TCPBUFSIZE	2063
TCPNODELAY	2064
TCPPOINT	2064
TCPWINDOWSIZE	2065
TECBEGINEVENTLOGGING	2065
TECHOST	2066
TECPOINT	2066
TECUTF8EVENT	2066
THROUGHPUTDATATHRESHOLD	2067
THROUGHPUTTIMETHRESHOLD	2067
TIMEFORMAT	2068
TXNGROUPMAX	2068
UNIQUETDPTECEVENTS	2069
UNIQUETECEVENTS	2070
USEREXIT	2070
VERBCHECK	2071
VOLUMEHISTORY	2071
Server utilities	2071
DSMMAXSG (Increase the block size for writing data)	2072
DSMSERV (Start the server)	2073
Server startup script: rc.dsmserv	2075
Server startup script: dsmserv.rc	2075
DSMSERV DISPLAY DBSPACE (Display information about database storage space)	2076
DSMSERV DISPLAY LOG (Display recovery log information)	2077
DSMSERV EXTEND DBSPACE (Increase space for the database)	2079
DSMSERV FORMAT (Format the database and log)	2080
DSMSERV INSERTDB (Move a server database into an empty database)	2082
DSMSERV LOADFORMAT (Format a database)	2084
DSMSERV REMOVEDB (Remove a database)	2085
DSMSERV RESTORE DB (Restore the database)	2087
DSMSERV RESTORE DB (Restore a database to its most current state)	2087
DSMSERV RESTORE DB (Restore a database to a point-in-time)	2089
DSMSERV UPDATE (Create registry entries for a server instance)	2093
DSMULOG (Capture IBM Spectrum Protect server messages to a user log file)	2093
Device utilities	2094
AIX and Linux: dsmsanlist (Display information about devices)	2094
Linux: autoconf (Auto configure devices)	2095
Windows: tsmdlst (Display information about devices)	2097
Server scripts and macros for automation	2098
Server scripts	2098
Defining a server script	2099
Running commands in parallel or serially	2099
Continuing commands across multiple command lines	2100

Including substitution variables in a script	2101
Including logic flow statements in a script	2101
Specifying the IF clause	2101
Specifying the EXIT statement	2102
Specifying the GOTO statement	2102
Using SELECT commands in a script	2102
Updating a script	2103
Appending a new command	2103
Replacing an existing command	2104
Adding a command and line number	2104
Deleting a command from a server script	2104
Querying a server script to create another server script	2104
Running a server script	2105
Administrative client macros	2105
Writing commands in a macro	2106
Writing comments in a macro	2106
Including continuation characters in a macro	2106
Including substitution variables in a macro	2107
Running a macro	2108
Command processing in a macro	2108
Return codes for use in IBM Spectrum Protect scripts	2109
PDF files	2111
Klienty	2111
Interfejs API	2111
Wydajność	2112
Rozwiązywanie problemów	2112
Messages, return codes, and error codes	2112
Introduction to messages	2112
IBM Spectrum Protect server and client messages format	2113
Interpreting return code messages	2114
Example one for QUERY EVENT command	2114
Example two for DEFINE VOLUME command	2114
ANE messages	2114
ANR messages	2115
ANS 0000-9999 messages	2115
API return codes	2115
I/O code descriptions in server messages	2115
Device drivers completion code and operation code descriptions overview	2116
Completion code values common to all device classes	2116
Completion code values for media changers	2117
Completion code values for tape drives	2119
Standard ASC and ASCQ codes descriptions	2120
Device error codes in the AIX system error log	2123
IBM Global Security Kit return codes	2124
Glosariusz	2133
A	2133
B	2134
C	2135

D	2135
E	2136
F	2136
G	2136
H	2137
I	2137
K	2137
L	2139
M	2140
N	2141
O	2142
P	2143
R	2146
S	2147
Ś	2149
T	2149
U	2150
V	2151
W	2151
Z	2152

Dokumentacja produktu IBM Spectrum Protect

Produkt IBM Spectrum Protect udostępnia zautomatyzowane, centralnie planowane, sterowane za pomocą strategii funkcje tworzenia kopii zapasowych i archiwalnych oraz zarządzania pamięcią masową dla serwerów plików, stacji roboczych, maszyn wirtualnych i aplikacji. Dokumentacja produktu IBM Spectrum Protect ułatwia przygotowanie, skonfigurowanie i zarządzanie rozwiązaniami do ochrony danych.

Pierwsze kroki

Instalowanie i aktualizowanie serwerów
Instalowanie i aktualizowanie Operations Center
Wybieranie i implementowanie rozwiązań ochrony danych
Nowości dla serwera
[🔗 Filmy wideo dla nowości](#)
Pliki PDF


Typowe czynności

Codziennie zadania monitorowania
Dodawanie klientów
Replikowanie danych klientów do innego serwera
Zarządzanie serwerem, klientami i Operations Center
Konfigurowanie pamięci masowej
Komendy, opcje i programy narzędziowe serwera

Rozwiązywanie problemów i wsparcie

Rozwiązywanie problemów
Optymalizowanie wydajności
[🔗 Najnowsze pakiety poprawek dla klientów i serwerów produktu IBM Spectrum Protect](#)
[🔗 Dział wsparcia IBM dla oprogramowania](#)

Więcej informacji

 Wskazówki dla użytkowników Centrum wiedzy IBM®
Pakiety produktów i produkty pokrewne
[🔗 Strona główna rodziny produktów](#)
[🔗 Serwis wiki dla produktów IBM Spectrum Protect](#)
[🔗 IBM Spectrum Protect Developer Center](#)
[🔗 Dokumentacja techniczna IBM Redbooks](#)
[🔗 IBM Skills Gateway for Systems](#)
Dostępność
Klauzula prawna

© Copyright IBM Corp. 1993, 2018

Ułatwienia dostępu w rodzinie produktów IBM Spectrum Protect

Funkcje dostępności ułatwiają posługiwanie się oprogramowaniem osobom niepełnosprawnym, np. z ograniczonymi możliwościami poruszania się lub z osłabionym wzrokiem.

Przegląd

Rodzina produktów IBM Spectrum Protect™ obsługuje następujące główne funkcje ułatwień dostępu:

- Możliwość obsługi tylko przy użyciu klawiatury
- Operacje używające oprogramowania lektora ekranowego

Rodzina produktów IBM Spectrum Protect używa najnowszego standardu W3C WAI-ARIA 1.0 do zachowania zgodności z przepisami US Section 508 i Web Content Accessibility Guidelines (WCAG) 2.0. Aby skorzystać z ułatwień dostępu, należy użyć najnowszej wersji lektora ekranowego w połączeniu z najnowszą przeglądarką WWW obsługiwana przez ten produkt.

Dokumentacja produktu w Centrum Wiedzy IBM obsługuje ułatwienia dostępu. Funkcje ułatwień dostępu w Centrum Wiedzy IBM zostały opisane w sekcji Ułatwienia dostępu w Centrum Wiedzy IBM.

Nawigowanie przy użyciu klawiatury

Ten produkt używa standardowych klawiszy nawigacyjnych.

Informacje dotyczące interfejsów

Interfejs użytkownika produktu nie stosuje elementów, które migają z częstotliwością od 2 do 55 razy na sekundę.

Internetowy interfejs użytkownika wykorzystuje arkusze stylów CSS do wyświetlania treści. Aplikacja pozwala użytkownikom z wadą wzroku na używanie innych ustawień wyświetlania, w tym trybu wyświetlania o wysokim kontraście. Użytkownik może kontrolować wielkość czcionki za pomocą urządzenia lub ustawień przeglądarki WWW.

Interfejsy użytkownika zawierają elementy nawigacyjne WAI-ARIA, które umożliwiają szybkie przechodzenie między obszarami funkcjonalnymi aplikacji.

Oprogramowanie dostawcy

Rodzina produktów IBM Spectrum Protect zawiera oprogramowanie pochodzące od innych dostawców, które nie jest objęte umową licencyjną firmy IBM. Firma IBM nie udziela jakichkolwiek gwarancji w zakresie ułatwień dostępu w tych produktach. Skontaktuj się z dostawcą, aby uzyskać informacje o ułatwieniach dostępu w tych produktach.

Informacje pokrewne na temat ułatwień dostępu

Oprócz standardowych serwisów WWW działów pomocy telefonicznej i wsparcia firma IBM wprowadziła usługę telefonu tekstowego (TTY) dla niesłyszących lub niedosłyszących klientów pozwalającą uzyskać dostęp do usług handlowych i wsparcia:

Usługa TTY
800-IBM-3383 (800-426-3383)
(w Ameryce Północnej)

Więcej informacji na temat zaangażowania firmy IBM w upowszechnianie ułatwień dostępu można znaleźć w sekcji IBM Accessibility .

Pakiety produktów i produkty pokrewne

IBM Spectrum Storage Suite i produkty pokrewne rozbudowują funkcje podstawowego produktu IBM Spectrum Protect.

Pakiety produktów i licencjonowanie

Produkty IBM Spectrum Protect i IBM Spectrum Protect Extended Edition oferują podstawowe komponenty pozwalające na automatyzację i centralne zarządzanie operacjami tworzenia i odtwarzania kopii zapasowych. Komponenty serwera oraz klienta kopii zapasowych i archiwalnych zawierają podstawowe funkcje, takie jak tworzenie i odtwarzanie kopii zapasowych oraz archiwizowanie i odtwarzanie plików, katalogów oraz obrazów dyskowych.

Dokumentacja produktu zawiera informacje dotyczące produktów IBM Spectrum Protect i IBM Spectrum Protect Extended Edition.

Pakiety produktów, które łączą IBM Spectrum Protect z produktami pokrewnymi, mogą umożliwiać prostsze kupowanie i zarządzanie oprogramowaniem. Pakiety zawierają produkty, które mogą spełniać szereg wymagań dotyczących ochrony danych i odtwarzania, udostępniając uproszczony model licencjonowania.

Wybierz zestaw produktów spełniających wymagania biznesowe:

- Informacje na temat zestawów produktów IBM Spectrum Protect zawiera nota techniczna 7048916.

- Informacje na temat produktu IBM Spectrum Storage Suite, który zawiera produkt IBM Spectrum Protect i inne produkty, podano w sekcji IBM Spectrum Storage Suite.

Produkty pokrewne

Produkt IBM Spectrum Protect można rozbudować, dodając funkcje dostępne w produktach pokrewnych.

Produkt	Kluczowe korzyści	Odsyłacze
IBM® Cloud Object Storage	Udostępnia platformę służącą do przechowywania nieustrukturyzowanych danych, z woluminami o wielkości od petabajtów do eksabajtów.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup
IBM Spectrum Control	Pozwala na zarządzanie danymi analitycznymi.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Copy Data Management	Ten produkt kataloguje obrazy stanu NetApp i VMware w celu ułatwienia zarządzaniem z wykorzystaniem ról i odtwarzaniem danych z kopii zapasowych.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect High Speed Data Transfer	Ten produkt pozwala na użycie technologii FASP (Fast Adaptive Secure Protocol) pozwalającej na zwiększenie szybkości przesyłania danych w środowisku, w którym wykryto problemy z wydajnością sieci WAN.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Określenie, czy technologia Aspera FASP może zoptymalizować przesyłanie danych w używanym środowisku
IBM Spectrum Protect for Data Retention	Zapewnia długoterminową ochronę w przypadku archiwizowania rekordów biznesowych, plików lub danych. Archiwizowanie danych zgodnie z obowiązującymi przepisami wymaga dodatkowych zabezpieczeń umożliwiających zapewnienie czasu przechowywania danych. Dzięki temu można uzyskać pewność, że dane nie zostaną usunięte przedwcześnie (nieumyślnie lub intencjonalnie). Aby spełnić wymagania, IBM Spectrum Protect for Data Retention zapewnia dodatkową ochronę danych archiwizowanych przez produkt IBM Spectrum Protect.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu <p>Wskazówka: Dokumentacja dla tego produktu jest zawarta w dokumentacji produktu IBM Spectrum Protect.</p>
IBM Spectrum Protect Plus	Zapewnia ochronę danych i wysoką dostępność dla środowisk wirtualnych. Można go wdrożyć w ciągu kilku minut i objąć środowisko ochroną w czasie poniżej godziny. Można użyć wersji autonomicznej produktu IBM Spectrum Protect Plus lub zintegrować go z używanym środowiskiem IBM Spectrum Protect w celu długoterminowego przechowywania danych i zapewnienia nadzoru w sposób skalowalny i wydajny.	

Produkt	Kluczowe korzyści	Odsyłacze
IBM Spectrum Protect Snapshot	<p>Produkt chroni dane dzięki zintegrowanej funkcji tworzenia i odtwarzania kopii zapasowej obrazu stanu (z rozpoznawaniem aplikacji).</p> <p>Dane przechowywane przez aplikacje IBM DB2, SAP, Oracle, Microsoft Exchange i Microsoft SQL Server mogą być chronione przez produkt IBM Spectrum Protect Snapshot. Dzięki temu można tworzyć i zarządzać obrazami stanu na poziomie woluminów dla systemów plików i aplikacji niestandardowych. Można określić, czy produkt IBM Spectrum Protect Snapshot ma być zintegrowany z produktem IBM Spectrum Protect.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect for Databases	<p>Chroni dane zapisane w bazach danych Oracle i Microsoft SQL dzięki zautomatyzowanym zadaniom, narzędziom i interfejsom. To oprogramowanie tworzy na bieżąco spójne i scentralizowane kopie zapasowe, chroniąc istotne dane, jednocześnie minimalizując koszty działania.</p> <p>Wskazówka: Obsługa otwartych baz danych IBM DB2 i IBM Informix jest zapewniana przez serwery IBM Spectrum Protect. Do tworzenia kopii zapasowych tych baz danych nie trzeba instalować produktu IBM Spectrum Protect for Databases. Więcej informacji zawiera dokumentacja produktów DB2 i Informix.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect for Enterprise Resource Planning	<p>Udostępnia ochronę danych systemu SAP.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect for Mail	<p>Automatyzuje ochronę danych, dzięki czemu kopie zapasowe są tworzone bez zamykania serwerów Microsoft Exchange lub IBM Domino.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect for Space Management	<p>Produkt hierarchicznego zarządzania pamięcią masową, który zmniejsza koszty pamięci masowej dla rzadko używanych informacji, bez modyfikowania aplikacji i zmiany sposobu pracy użytkowników. Ten produkt jest dostępny dla systemów operacyjnych IBM AIX i Linux.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Protect HSM for Windows	<p>Produkt hierarchicznego zarządzania pamięcią masową, który zmniejsza koszty pamięci masowej dla rzadko używanych informacji, bez modyfikowania aplikacji i zmiany sposobu pracy użytkowników. Ten produkt jest dostępny dla systemu operacyjnego Windows.</p>	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu

Produkt	Kluczowe korzyści	Odsyłacze
IBM Spectrum Protect for SAN	Współpracuje z serwerami i komputerami klienckimi, przesyłając dane w sieci pamięci masowej (SAN) zamiast w sieci LAN. Ten produkt jest agentem pamięci masowej, który umożliwia tworzenie i odtwarzanie kopii zapasowych bez obciążania sieci LAN.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu <p>Wersja dokumentacji produktu: Dokumentacja produktu IBM Tivoli Storage Manager for SAN 7.1 jest również aktualna dla rodziny produktów IBM Spectrum Protect 8.1.</p>
IBM Spectrum Protect for Virtual Environments	Zapewnia ochronę dostosowaną do wirtualnych środowisk VMware i Hyper-V.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Spectrum Scale	Udostępnia skalowalną pamięć masową dla nieustrukturyzowanych danych.	<ul style="list-style-type: none"> • Dowiedz się więcej i kup • Dokumentacja produktu
IBM Tivoli Storage Manager for z/OS Media	Zarządza zasobami dyskowymi i taśmowymi w systemie z/OS dla serwerów IBM Spectrum Protect działających w systemach AIX lub Linux na platformie System z.	<ul style="list-style-type: none"> • Dokumentacja produktu

Pliki PDF

Gotowe pliki PDF można pobrać z Centrum Wiedzy IBM® lub z serwisu FTP.

Gotowe pliki PDF

W podanych niżej sekcjach zamieszczono pliki PDF dostępne dla tej wersji:

- Rozwiązania ochrony danych
- Serwery

Pakiet plików PDF

Pobierz pakiet, który zawiera wszystkie pliki PDF dla tego wydania:

<ftp://public.dhe.ibm.com/software/products/ISP/current/>

Aktualizacje w tej wersji

Zapoznaj się z nowymi funkcjami i udoskonaleniami produktów, aby określić potencjalne korzyści dla operacji zarządzania pamięcią masową. Uwagi do wydania zawierają odsyłacze, do których można uzyskać dostęp w celu uzyskania ważnych informacji przed rozpoczęciem instalacji lub aktualizacji produktów i komponentów.

Komponent	Podsumowanie aktualizacji	Uwagi do wydania dla wersji 8.1
Komponenty serwera	Aktualizacje	Uwagi do wydania

- Program Beta
Program Beta produktu IBM Spectrum Protect pozwala na wgląd w funkcje tworzonego produktu i umożliwia przesłanie opinii mających wpływ na projekt produktu. Dzięki temu programowi można przetestować produkt we własnym środowisku i przekazać uwagi bezpośrednio do programistów.

IBM Spectrum Protect concepts

IBM Spectrum Protect™ provides a comprehensive data protection environment.

- IBM Spectrum Protect overview
IBM Spectrum Protect provides centralized, automated data protection that helps to reduce data loss and manage

compliance with data retention and availability requirements.

- Data storage concepts in IBM Spectrum Protect
IBM Spectrum Protect provides functions to store data in a range of device and media storage.
- Data protection strategies with IBM Spectrum Protect
IBM Spectrum Protect provides ways for you to implement various data protection strategies.

IBM Spectrum Protect overview

IBM Spectrum Protect™ provides centralized, automated data protection that helps to reduce data loss and manage compliance with data retention and availability requirements.

- Data protection components
The data protection solutions that IBM Spectrum Protect provides consist of a server, client systems and applications, and storage media. IBM Spectrum Protect provides management interfaces for monitoring and reporting the data protection status.
- Data protection services
IBM Spectrum Protect provides data protection services to store and recover data from various types of clients. The data protection services are implemented through policies that are defined on the server. You can use client scheduling to automate the data protection services.
- Processes for managing data protection with IBM Spectrum Protect
The IBM Spectrum Protect server inventory has a key role in the processes for data protection. You define policies that the server uses to manage data storage.
- User interfaces for the IBM Spectrum Protect environment
For monitoring and configuration tasks, IBM Spectrum Protect provides various interfaces, including the Operations Center, a command-line interface, and an SQL administrative interface.

Data protection components

The data protection solutions that IBM Spectrum Protect™ provides consist of a server, client systems and applications, and storage media. IBM Spectrum Protect provides management interfaces for monitoring and reporting the data protection status.

Server

Client systems send data to the server to be stored as backups or archived data. The server includes an *inventory*, which is a repository of information about client data.

The inventory includes the following components:

Database

Information about each file, logical volume, or database that the server backs up, archives, or migrates is stored in the server database. The server database also contains information about the policy and schedules for data protection services.


Recovery log

Records of database transactions are kept in this log. The database uses the recovery log to ensure data consistency in the database.

Client systems and applications

Clients are applications, virtual machines, and systems that must be protected. The clients send data to the server, as shown in Figure 1.

Figure 1. Components in the data protection solution



Client software

For IBM Spectrum Protect to protect client data, the appropriate software must be installed on the client system and the client must be registered with the server.

Client nodes

A *client node* is equivalent to a computer, virtual machine, or application, such as a backup-archive client that is installed on a workstation for file system backups. Each client node must be registered with the server. Multiple nodes can be registered on a single computer.

Storage media

The server stores client data to storage media. The following types of media are used:

Storage devices

The server can write data to hard disk drives, disk arrays and subsystems, stand-alone tape drives, tape libraries, and other types of random-access and sequential-access storage. Storage devices can be connected directly to the server or connected through a local area network (LAN) or a storage area network (SAN).

Storage pools

Storage devices that are connected to the server are grouped into *storage pools*. Each storage pool represents a set of storage devices of the same media type, such as disk or tape drives. IBM Spectrum Protect stores all of the client data in storage pools. You can organize storage pools into a *hierarchy*, so that data storage can transfer from disk storage to lower-cost storage such as tape devices.


Data protection services

IBM Spectrum Protect™ provides data protection services to store and recover data from various types of clients. The data protection services are implemented through policies that are defined on the server. You can use client scheduling to automate the data protection services.

Types of data protection services

IBM Spectrum Protect provides services to store and recover client data as shown in Figure 1.

Figure 1. Data protection services



IBM Spectrum Protect provides the following types of data protection services:

Back up and restore services

You run a backup process to create a copy of a *data object* that can be used for recovery if the original data object is lost. A data object can be a file, a directory, or a user-defined data object, such as a database.

To minimize the use of system resources during the backup operation, IBM Spectrum Protect uses the *progressive incremental backup* method. For this backup method, a first full backup of all data objects is created and in subsequent backup operations only changed data is moved to storage. Compared to incremental and differential backup methods that require taking periodic full backups, the progressive incremental backup method provides the following benefits:

- Reduces data redundancy
- Uses less network bandwidth
- Requires less storage pool space

To further reduce storage capacity requirements and network bandwidth usage, IBM Spectrum Protect includes *data deduplication* for data backups. The data deduplication technique removes duplicate data extents from backups.

You run a restore process to copy an object from a storage pool to the client. You can restore a single file, all files in a directory, or all of the data on a computer.

Archive and retrieve services

You use the archive service to preserve data that must be stored for a long time, such as for regulatory compliance. The archive service provides the following features:

- When you archive data, you specify how long the data must be stored.
- You can request that files and directories are copied to long-term storage on media. For example, you might choose to store this data on a tape device, which can reduce the cost of storage.
- You can specify that the original files are erased from the client after the files are archived.

The retrieve service provides the following features:

- When you retrieve data, the data is copied from a storage pool to a client node.
- The retrieve operation does not affect the archive copy in the storage pool.

Migrate and recall services

You use migrate and recall services to manage space on client systems. The goal of space management is to maximize available media capacity for new data and to minimize access time to data. You can migrate data to server storage to maintain sufficient free storage space on a local file system. You can store migrated data in the following ways:

- On disk storage for long-term storage
- In a *virtual tape library* (VTL) for fast recall of files

You can recall files to the client node on demand, either automatically or selectively.

Types of client data that can be protected

You can protect data for the following types of clients with IBM Spectrum Protect:

Application clients

IBM Spectrum Protect can protect data for specific products or applications. These clients are called *application clients*. To protect the *structured data* for these clients, in other words the data in database fields, you must back up components that are specific to the application. IBM Spectrum Protect can protect the following applications:

- IBM Spectrum Protect for Enterprise Resource Planning clients:
 - Data Protection for SAP HANA
 - Data Protection for SAP for DB2®
 - Data Protection for SAP for Oracle
- IBM Spectrum Protect for Databases clients:
 - Data Protection for Microsoft SQL server
 - Data Protection for Oracle
- IBM Spectrum Protect for Mail clients:
 - Data Protection for IBM® Domino®
 - Data Protection for Microsoft Exchange Server

Virtual machines

Virtual machines that are backed up by using application client software that is installed on the virtual machine. In the IBM Spectrum Protect environment, a virtual machine can be protected by the IBM Spectrum Protect for Virtual Environments. System clients

The following IBM Spectrum Protect clients are called *system clients*:

- All clients that back up data in files and directories, in other words *unstructured data*, such as backup-archive clients and API clients that are installed on workstations.
- A server that is included in a server-to-server virtual volume configuration.
- A virtual machine that is backed up by using backup-archive client software that is installed on the virtual machine.


Processes for managing data protection with IBM Spectrum Protect

The IBM Spectrum Protect™ server inventory has a key role in the processes for data protection. You define policies that the server uses to manage data storage.

Data management process

Figure 1 shows the IBM Spectrum Protect data management process.

Figure 1. Data management process



IBM Spectrum Protect uses policies to control how the server stores and manages data objects on various types of storage devices and media. You associate a client with a policy domain that contains one active policy set. When a client backs up, archives, or migrates a file, the file is bound to a management class in the active policy set of the policy domain. The management class and the backup and archive copy groups specify where files are stored and how they are managed. If you set up server storage in a hierarchy, you can migrate files to different storage pools.

Inventory components

The following inventory components are key to the operation of the server:

Server database

The server database contains information about client data and server operations. The database stores information about client data, called *metadata*. Information about client data includes the file name, file size, file owner, management class, copy group, and location of the file in server storage. The database includes the following information that is necessary for the operation of the server:

- Definitions of client nodes and administrators
- Policies and schedules
- Server settings
- Records of server operations, such as activity logs and event records
- Intermediate results for administrative queries

Recovery log

The server records database transactions in the recovery log. The recovery log helps to ensure that a failure does not leave the database in an inconsistent state. The recovery log is also used to maintain consistency across server start operations. The recovery log consists of the following logs:

Active log

This log records current transactions on the server. This information is required to start the server and database after a disaster.

Log mirror (optional)

The active log mirror is a copy of the active log that can be used if the active log files cannot be read. All changes that are made to the active log are also written to a log mirror. You can set up one active log mirror.

Archive log

The archive log contains copies of closed log files that were in the active log. The archive log is included in database backups and is used for recovery of the server database. Archive log files that are included in a database backup are automatically pruned after a full database backup cycle is complete. The archive log must have enough space to store the log files for database backups.

Archive failover log (optional)

The archive failover log, also called a secondary archive log, is the directory that the server uses to store archive log files when the archive log directory is full.

Policy-based data management

In the IBM Spectrum Protect environment, a *policy* for data protection management contains rules that determine how client data is stored and managed. The primary purpose of a policy is to implement the following data management objectives:

- Control which storage pool client data is initially stored in
- Define retention criteria that controls how many copies of objects are stored
- Define how long copies of objects are retained

Policy-based data management helps you to focus on the business requirements for protecting data rather than on managing storage devices and media. Administrators define policies and assign client nodes to a *policy domain*.

Depending on your business needs, you can have one policy or many. In a business organization, for example, different departments with different types of data can have customized storage management plans. Policies can be updated, and the updates can be applied to data that is already managed.

When you install IBM Spectrum Protect, a default policy that is named STANDARD is already defined. The STANDARD policy provides basic backup protection for user workstations. To provide different levels of service for different clients, you can add to the default policy or create a new policy.

You create policies by defining the following policy components:

Policy domain

The policy domain is the primary organizational method of grouping client nodes that share common rules for data management. Although a client node can be defined to more than one server, the client node can be defined to only one policy domain on each server.

Policy set

A *policy set* is a number of policies that are grouped so that the policy for the client nodes in the domain can be activated or deactivated as required. An administrator uses a policy set to implement different management classes based on business and user needs. A policy domain can contain multiple policy sets, but only one policy set can be active in the domain. Each policy set contains a default management class and any number of extra management classes.

Management class

A *management class* is a policy object that you can bind to each category of data to specify how the server manages the data. There can be one or more management classes. One management class is assigned to be the default management class that is used by clients unless they specifically override the default to use a specific management class.

The management class can contain a backup copy group, an archive copy group, and space management attributes. A copy group determines how the server manages backup versions or archived copies of the file. The space management attributes determine whether the file is eligible for migration by the space manager client to server storage, and under what conditions the file is migrated.

Copy group

A *copy group* is a set of attributes in a management class that controls the following factors:

- Where the server stores versions of backed up files or archive copies
- How long the server keeps versions of backed up files or archive copies
- How many versions of backup copies are retained
- What method to use to generate versions of backed up files or archive copies

Security management

IBM Spectrum Protect includes security features for registration of administrators and users. After administrators are registered, they must be granted authority by being assigned one or more administrative privilege classes. An administrator with system privilege can perform any server function. Administrators with policy, storage, operator, or node privileges can perform subsets of server functions. The server can be accessed by using the following methods, each controlled with a password:

- Administrator access to manage the server
- Client access to nodes to store and retrieve data

Also included are features that can help to ensure security when clients connect to the server. Depending on business requirements, as an administrator, you can choose one of the following client registration methods:

Open registration

When the client first connects to the server, the user is requested for a node name, password, and contact information.

Open registration provides the user with following default settings:

- The client node is assigned to the STANDARD policy domain.
- The user can define whether files are compressed to decrease the amount of data that is sent over networks and the space that is occupied by the data in storage.
- The user can delete archived copies of files from server storage, but not backup versions of files.

Closed registration

Closed registration is the default method for client registration to the server. For this type of registration, an administrator registers all clients. The administrator can implement the following settings:

- Assign the node to any policy domain
- Determine whether the user can use compression or not, or if the user can choose
- Control whether the user can delete backed up files or archived files

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL). SSL is the standard technology that you use to create encrypted sessions for servers and clients, and provides a secure channel to communicate over open communication paths. With SSL, the identity of the server is verified by using digital certificates. If you authenticate with a Lightweight Directory Access Protocol (LDAP) server, passwords between the server and the LDAP server are protected by Transport Layer Security (TLS). The TLS protocol is the successor to the SSL protocol. When a server and client communicate, TLS ensures that third parties cannot intercept messages.


User interfaces for the IBM Spectrum Protect environment

For monitoring and configuration tasks, IBM Spectrum Protect™ provides various interfaces, including the Operations Center, a command-line interface, and an SQL administrative interface.

Interfaces for data storage management

The Operations Center is the primary interface for administrators to monitor and administer servers. A key benefit of the Operations Center is that you can monitor multiple servers, as shown in Figure 1. You can also monitor and administer IBM Spectrum Protect from a command-line administrative interface.

Figure 1. User interfaces for data storage management



You use the following interfaces to interact with IBM Spectrum Protect:

Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment. You can use the Operations Center to complete monitoring and certain administration tasks, for example:

- You can monitor multiple servers and clients.
- You can monitor the transaction activity for specific components in the data path, such as the server database, the recovery log, storage devices, and storage pools.

Command-line interface

You can use a command-line interface to run administration tasks for servers. You can access the command-line interface through either the IBM Spectrum Protect administrative client or the Operations Center.

Access to information in the server database by using SQL statements

You can use SQL SELECT statements to query the server database and display the results. Third-party SQL tools are available to aid administrators in database management.

Interfaces for client activity management

IBM Spectrum Protect provides the following types of interfaces for managing client activity:

- An application programming interface (API)
- Graphical user interfaces for clients
- Browser interface for the backup-archive client
- Command-line interfaces for clients

Data storage concepts in IBM Spectrum Protect

IBM Spectrum Protect™ provides functions to store data in a range of device and media storage.

To make storage devices available to the server, you must attach the storage devices and map storage pools to device classes, libraries, and drives.

- Types of storage devices
You can use various storage devices with IBM Spectrum Protect to meet specific data protection goals.
- Data storage in storage pools
Logical storage pools are the principal components in the IBM Spectrum Protect model of data storage. You can optimize the usage of storage devices by manipulating the properties of storage pools and volumes.
- Data transport to storage across networks
The IBM Spectrum Protect environment provides ways to securely move data to storage across various types of networks

and configurations.

Types of storage devices

You can use various storage devices with IBM Spectrum Protect™ to meet specific data protection goals.

Storage devices and storage objects

The IBM Spectrum Protect server can connect to a combination of manual and automated storage devices. You can connect the following types of storage devices to IBM Spectrum Protect:

- Disk devices that are directly attached, SAN-attached, or network attached
- Physical tape devices that are either manually operated or automated
- Virtual tape devices
- Cloud object storage

IBM Spectrum Protect represents physical storage devices and media with storage objects that you define in the server database. Storage objects classify available storage resources and manage migration from one storage pool to another. Table 1 describes the storage objects in the server storage environment.

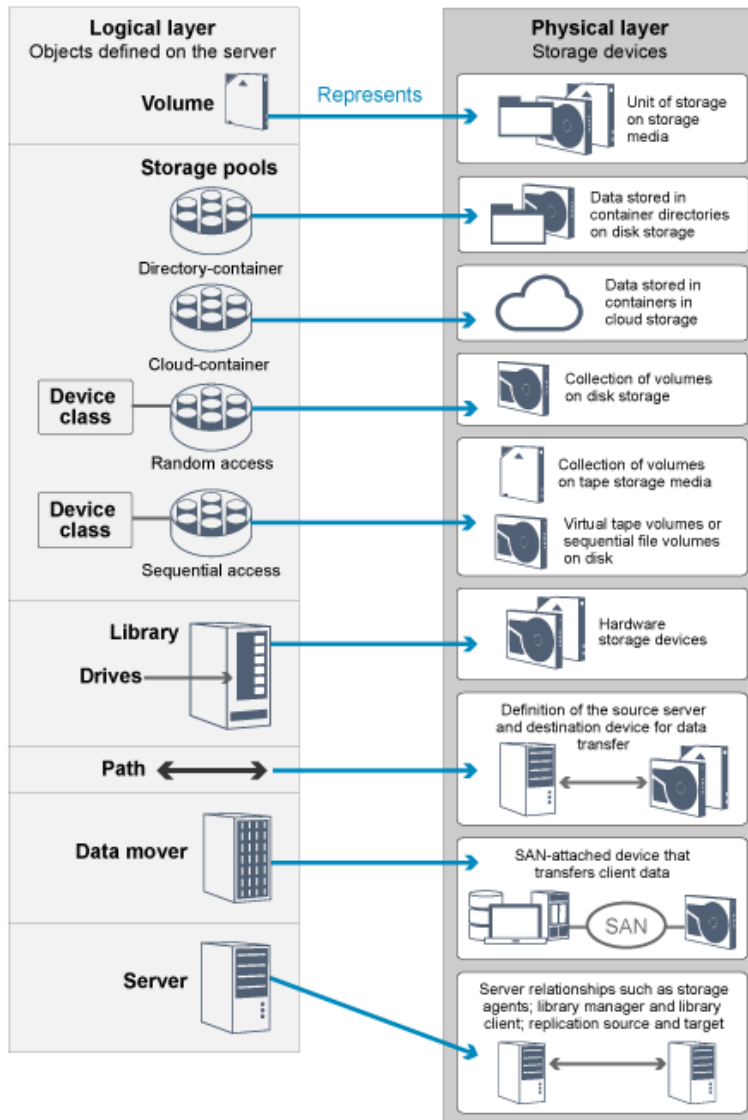
Table 1. Storage objects and representations

Storage object	What the object represents
Volume	A discrete unit of storage on disk, tape, or other storage media. Each volume is associated with a single storage pool.
Storage pool	A set of storage volumes or containers that is the destination that is used to store client data. IBM Spectrum Protect uses the following types of storage pool: <ul style="list-style-type: none">• Directory-container storage pools• Cloud-container storage pools• Sequential-access storage pools that are associated with a device class• Random-access storage pools that are associated with a device class
Container	A data storage location, for example, a file, directory, or device.
Container storage pool	A primary storage pool that a server uses to store data. Data is stored in containers in file system directories or in cloud storage. Data is deduplicated, if necessary, as the server writes data to the storage pool.
Device class	The type of storage device that can use the volumes that are defined in a sequential-access or random-access storage pool. Each device class of removable media type is associated with a single library.
Library	A storage device. For example, a library can represent a stand-alone drive, a set of stand-alone drives, a multiple-drive automated device, or a set of drives that is controlled by a media manager.
Drive	An object of a tape library device that provides the capability to read and write data to tape library media. Each drive is associated with a single library.
Path	The specification of the data source and the device destination. Before a storage device can be used, a path must be defined between the device and the source server that is moving data.

Storage object	What the object represents
Data mover	A SAN-attached device that is used to transfer client data. A data mover is used only in a data transfer where the server is not present, such as in a Network Data Management Protocol (NDMP) environment. Data movers transfer data between storage devices without using significant server, client, or network resources.
Server	A server that is managed by another IBM Spectrum Protect server.

The administrator defines the storage objects in the logical layer of the server, as illustrated in Figure 1.

Figure 1. Storage objects



Disk devices

You can store client data on disk devices with the following types of volumes:

- Directories in directory-container storage pools
- Random-access volumes of device type DISK
- Sequential-access volumes of device type FILE

IBM Spectrum Protect offers the following features when you use directory-container storage pools for data storage:

- You can apply data deduplication and disk caching techniques to maximize data storage usage.

- You can retrieve data from disk much faster than you can retrieve data from tape storage.

Physical tape devices

In a physical tape library, the storage capacity is defined in terms of the total number of volumes in the library. Physical tape devices can be used for the following activities:

- Storing client data that is backed up, archived, or migrated from client nodes
- Storing database backups
- Exporting data to another server or offsite storage

Moving data to tape provides the following benefits:

- You can keep data for clients on a disk device at the same time that the data is moved to tape.
- You can improve tape drive performance by streaming the data migration from disk to tape.
- You can spread out the times when the drives are in use to improve the efficiency of the tape drives.
- You can move data on tape to off-site vaults.
- You can limit power consumption because tape devices do not consume power after data is written to tape.
- You can apply encryption that is provided by the tape drive hardware to protect the data on tape.

Compared to equivalent disk and virtual tape storage, the unit cost to store data tends to be much less for physical tape devices.

Virtual tape libraries

A virtual tape library (VTL) does not use physical tape media. When you use VTL storage, you emulate the access mechanisms of tape hardware. In a VTL, you can define volumes and drives to provide greater flexibility for the storage environment. The storage capacity of a VTL is defined in terms of total available disk space. You can increase or decrease the number and size of volumes on disk.

Defining a VTL to the IBM Spectrum Protect server can improve performance because the server handles mount point processing for VTLs differently than for real tape libraries. Although the logical limitations of tape devices are still present, the physical limitations for tape hardware are not applicable to a VTL thus affording better scalability. You can use the IBM Spectrum Protect VTL when the following conditions are met:

- Only one type and generation of drive and media is emulated in the VTL.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

Data storage in storage pools

Logical storage pools are the principal components in the IBM Spectrum Protect™ model of data storage. You can optimize the usage of storage devices by manipulating the properties of storage pools and volumes.

Types of storage pools

The group of storage pools that you set up for the server is called *server storage*. You can define the following types of storage pools in server storage:

Primary storage pools

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files that are migrated from client nodes.

Copy storage pools

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class for space-managed files.

Container-copy storage pools

A named set of volumes that contain a copy of data extents that reside in directory-container storage pools. Container-copy storage pools are used only to protect the data that is stored in directory-container storage pools.


Active-data storage pools

A named set of storage pool volumes that contain only active versions of client backup data.

Primary storage pools

When you restore, retrieve, recall, or export file data, the requested file is obtained from a primary storage pool. Depending on the type of primary storage pool, the storage pools can be onsite or offsite. You can arrange primary storage pools in a storage hierarchy so that data can be transferred from disk storage to lower-cost storage such as tape devices. Figure 1 illustrates the concept of primary storage pools.

Figure 1. Primary storage pools



You can define the following types of primary storage pool:

Directory-container storage pools

A storage pool that the server uses to store data in containers in storage pool directories. Data that is stored in a directory-container storage pool can use either inline data deduplication, client-side data deduplication, inline compression, or client-side compression. Inline data deduplication or inline compression reduces data at the time it is stored.

Tip: Data that is compressed first cannot be deduplicated, however, deduplicated data can be compressed.

By using directory-container storage pools, you remove the need for volume reclamation, which improves server performance and reduces the cost of storage hardware. You can protect and repair data in directory-container storage pools at the level of the storage pool. You can tier data that is stored in a directory-container storage pool to a cloud-container storage pool.

Restriction: You cannot use any of the following functions with directory-container storage pools:

- Migration
- Reclamation
- Aggregation
- Collocation
- Simultaneous-write
- Storage pool backup
- Virtual volumes

Cloud-container storage pools

A storage pool that a server uses to store data in cloud storage. The cloud storage can be on premises or off premises. The cloud-container storage pools that are provided by IBM Spectrum Protect can store data to cloud storage that is object-based. By storing data in cloud-container storage pools, you can exploit the cost per unit advantages that clouds offer along with the scaling capabilities that cloud storage provides. You can use cloud tiering to lower costs by moving data from disk storage to a cloud-container storage pool. IBM Spectrum Protect manages the credentials, security, read and write I/Os,

and the lifecycle for data that is stored to the cloud. When cloud-container storage pools are implemented on the server, you can write directly to the cloud by configuring a cloud-container storage pool with the cloud credentials. Data that is stored in a cloud-container storage pool uses both inline data deduplication and inline compression. The server writes deduplicated, compressed, and encrypted data directly to the cloud. You can back up and restore data or archive and retrieve data directly from the cloud-container storage pool.

You can define the following types of cloud-container storage pools:

On premises

You can use the on premises type of cloud-container storage pool to store data in a private cloud, for more security and maximum control over your data. The disadvantages of a private cloud are higher costs due to hardware requirements and onsite maintenance.

Off premises

You can use the off premises type of cloud-container storage pool to store data in a public cloud. The advantage of using a public cloud is that you can achieve lower costs than for a private cloud, for example by eliminating maintenance. However, you must balance this benefit against possible performance issues due to connection speeds and reduced control over your data.

Storage pools that are associated with device classes

You can define a primary storage pool to use the following types of storage devices:

DISK device class

In a DISK device type of storage pool, data is stored in random access disk blocks. You can use caching in DISK storage pools to increase client restore performance with some limitations on server processing. Space allocation and tracking by blocks uses more database storage space and requires more processing power than allocation and tracking by volume.

FILE device class

In a FILE device type of storage pool, files are stored in sequential volumes for better sequential performance than for storage in disk blocks. To the server, these files have the characteristics of a tape volume so that this type of storage pool is better suited for migration to tape. FILE volumes are useful for *electronic vaulting*, where data is transferred electronically to a remote site rather than by physical shipment of tape. In general, this type of storage pool is preferred over DISK storage pools.

The server uses the following default random-access primary storage pools:

ARCHIVEPOOL

In the STANDARD policy, this storage pool is the destination for files that are archived from client nodes.

BACKUPPOOL

In the STANDARD policy, this storage pool is the destination for files that are backed up from client nodes.

SPACEMGPOOL


This storage pool is for space-managed files that are migrated from IBM Spectrum Protect for Space Management client nodes.

Copy storage pools

Copy storage pools contain active and inactive versions of data that is backed up from primary storage pools. A directory-container storage pool cannot be used as a copy storage pool. In addition, data from a directory-container storage pool cannot be copied into a copy storage pool. To protect directory-container storage pools, copy the data to a container-copy storage pool.

Figure 2 illustrates the concept of copy storage pools.

Figure 2. Copy storage pools



Copy storage pools provide a means of recovering from disasters or media failures. For example, when a client attempts to retrieve a damaged file from the primary storage pool, and the storage pool is unavailable or the file in the storage pool is corrupted, the client can restore the data from the copy storage pool.

You can move the volumes of copy storage pools offsite and still have the server track the volumes. Moving these volumes offsite provides a means of recovering from an onsite disaster. A copy storage pool can use sequential-access storage only, such as a tape device class or FILE device class.

Container-copy storage pools


A server can protect a directory-container storage pool by storing copies of the data in a container-copy storage pool. Data in container-copy storage pools is stored on tape volumes, which can be stored onsite or offsite. Damaged data in directory-container storage pools can be repaired by using deduplicated extents in container-copy storage pools. Container-copy storage pools provide an alternative to using a replication server to protect data in a directory-container storage pool.

Restriction: If all server data is lost, container-copy storage pools alone do not provide the same level of protection as replication:

- With replication, you can directly restore client data from the target server if the source server is unavailable.
- With container-copy storage pools, you must first restore the server from a database backup and then repair directory-container storage pools from tape volumes.

Figure 3 illustrates the concept of container-copy storage pools.

Figure 3. Container-copy storage pools



Depending on your system configuration, you can create protection schedules to simultaneously copy the directory-container storage pool data to onsite or offsite container-copy storage pools to meet your requirements:

- If replication is enabled, you can create one offsite container-copy pool. The offsite copy can be used to provide extra protection in a replicated environment.
- If replication is not enabled, you can create one onsite and one offsite container-copy storage pool.

Depending on the resources and requirements of your site, the ability to copy directory-container storage pools to tape has the following benefits:

- You avoid maintaining another server and more disk storage space.
- Data is copied to storage pools that are defined on the server. Performance is not dependent on, or affected by, the network connection between servers.
- You can satisfy regulatory and business requirements for offsite tape copies.


Active-data storage pools

An active-data pool contains only active versions of client backup data. In this case, the server does not have to position past inactive files that do not have to be restored. A directory-container storage pool cannot be used as an active-data storage pool. You use active-data pools to improve the efficiency of data storage and restore operations. For example, this type of storage pool can help you to achieve the following objectives:

- Increase the speed of client data restore operations
- Reduce the number of onsite or offsite storage volumes
- Reduce the amount of data that is transferred when you copy or restore files that are vaulted electronically in a remote location

Data that is migrated by hierarchical storage management (HSM) clients and archive data are not permitted in active-data pools. As updated versions of backup data are stored in active-data pools, older versions are removed as the remaining data is consolidated from many sequential-access volumes onto fewer, new sequential-access volumes. Figure 4 illustrates the concept of active-data storage pools.

Figure 4. Active-data storage pools



Active-data pools can use any type of sequential-access storage. However, the benefits of an active-data pool depend on the device type that is associated with the pool. For example, active-data pools that are associated with a FILE device class are ideal for fast client restore operations because of the following reasons:

- FILE volumes do not have to be physically mounted
- Client sessions that are restoring from FILE volumes in an active-data pool can access the volumes concurrently, which improves restore performance

Related information:

[Directory-container storage pools FAQs](#)

[Cloud-container storage pools FAQs](#)

Data transport to storage across networks

The IBM Spectrum Protect™ environment provides ways to securely move data to storage across various types of networks and configurations.


Network configurations for storage devices

IBM Spectrum Protect provides methods for configuring clients and servers on a local area network (LAN), on a storage area network (SAN), LAN-free data movement, and as network-attached storage.

Data backup operations over a LAN

Figure 1 shows the data path for IBM Spectrum Protect backup operations over a LAN.

Figure 1. IBM Spectrum Protect backup operations over a LAN




In a LAN configuration, one or more tape libraries are associated with a single IBM Spectrum Protect server. In this type of configuration, client data, electronic mail, terminal connection, application program, and device control information must all be handled by the same network. Device control information and client backup and restore data flow across the LAN.

Data backup operations over a SAN

Figure 2 shows the data path for IBM Spectrum Protect backup operations over a SAN.

Figure 2. IBM Spectrum Protect backup operations over a SAN



A SAN is a dedicated storage network that can improve system performance. On a SAN, you can consolidate storage and relieve the distance, scalability, and bandwidth limitations of LANs and wide area networks (WANs). By using IBM Spectrum Protect in a SAN, you can take advantage of the following functions:

- Share storage devices among multiple IBM Spectrum Protect servers. Devices that use the GENERICTAPE device type are not included.
- Move data from a client system directly to storage devices without using the LAN. LAN-free data movement requires the installation of a storage agent on the client system. The storage agent is available with the IBM Spectrum Protect for SAN product.

Through the storage agent, the client can directly back up and restore data to a tape library or shared file system such as GPFS™. The IBM Spectrum Protect server maintains the server database and recovery log, and acts as the library manager to control device operations. The storage agent on the client handles the data transfer to the device on the SAN. This implementation frees bandwidth on the LAN that would otherwise be used for client data movement.

- Share tape drives and libraries that are supported by the IBM Spectrum Protect server.
- Consolidate multiple clients under a single client node name in a General Parallel File System (GPFS) cluster.

Network-attached storage

Network-attached storage (NAS) file servers are dedicated storage servers whose operating systems are optimized for file-serving functions. NAS file servers typically interact with IBM Spectrum Protect through industry-standard network protocols, such as network data management protocol (NDMP) or as primary storage for random-access or sequential access storage pools. IBM Spectrum Protect provides the following basic types of configurations that use NDMP for backing up and managing NAS file servers:

- IBM Spectrum Protect backs up a NAS file server to a library device that is directly attached to the NAS file server. The NAS file server, which can be remote from the IBM Spectrum Protect server, transfers backup data directly to a drive in a SCSI-attached tape library. Data is stored in NDMP-formatted storage pools, which can be backed up to storage media that can be moved offsite for protection in case of an onsite disaster.
- IBM Spectrum Protect backs up a NAS file server over the LAN to a storage-pool hierarchy. In this type of configuration, you can store NAS data directly to disk, either random access or sequential access, and then migrate the data to tape. You can also use this type of configuration for system replication. Data can also be backed up to storage media that can be moved offsite. The advantage of this type of configuration is that you have all of the data management features associated with a storage pool hierarchy.
- The IBM Spectrum Protect client reads the data from the NAS system by using NFS or CIFS protocols and sends the data to the server to be stored.

Storage management

You manage the devices and media that are used to store client data through the IBM Spectrum Protect server. The server integrates storage management with the policies that you define for managing client data in the following areas:

Types of devices for server storage

With IBM Spectrum Protect, you can use directly attached devices and network-attached devices for server storage. IBM Spectrum Protect represents physical storage devices and media with administrator-defined storage objects.

Data migration through the storage hierarchy

For primary storage pools other than directory-container storage pools, you can organize the storage pools into one or more hierarchical structures. This storage hierarchy provides flexibility in a number of ways. For example, you can set a policy to back up data to disks for faster backup operations. The IBM Spectrum Protect server can then automatically migrate data from disk to tape.

Removal of expired data

The policy that you define controls when client data automatically expires from the IBM Spectrum Protect server. To remove data that is eligible for expiration, a server expiration process marks data as expired and deletes metadata for the expired data from the database. The space that is occupied by the expired data is then available for new data. You can control the frequency of the expiration process by using a server option.

Media reuse by reclamation

As server policies automatically expire data, the media where the data is stored accumulates unused space. For storage media other than directory-container storage pools or random disk storage pools, the IBM Spectrum Protect server implements *reclamation*, a process that frees media for reuse without traditional tape rotation. Reclamation automatically defragments media by consolidating unexpired data onto other media when the free space on media reaches a defined level. The reclaimed media can then be used again by the server. Reclamation allows media to be automatically circulated through the storage management process and minimize the number of media that are required.

Consolidating backed up client data

By grouping the client data that is backed up, you can minimize the number of media mounts required for client recovery. The IBM Spectrum Protect server provides the following methods for grouping client files on storage media other than directory-container storage pools:

Collocating client data

The IBM Spectrum Protect server can *collocate* client data, in other words store client data on a few volumes instead of spreading the data across many volumes. Collocation by client minimizes the number of volumes that are required to back up and restore client data. Data collocation might increase the number of volume mounts because each client might have a dedicated volume instead of data storage from several clients in the same volume.

You can set the server to collocate client data when the data is initially placed in server storage. In a storage hierarchy, you can collocate the data when the server migrates the data from the initial storage pool to the next storage pool in the storage hierarchy. You can collocate by client, by file space per client, or by a group of clients. Your selection depends on the size of the file spaces that are stored and restore requirements.

Associating active-data pools with various devices

Active-data pools are useful for fast restoration of client data. Benefits include a reduction in the number of onsite or offsite storage volumes, or reducing bandwidth when you copy or restore files that are vaulted electronically in a remote location. Active-data pools that use removable media, such as tape, offer similar benefits. Although tape devices must be mounted, the server does not have to position past inactive files. However, the primary benefit of using removable media in active-data pools is that the number of volumes that are used for onsite and offsite storage is reduced. If you store data to a remote location, you can minimize the amount of data that must be transferred by copying and restoring only active data.

Creating a backup set

A backup set contains all of the active backed-up files that exist for that client in server storage. The backup set is portable and is retained for the time that you specify. A backup set is in addition to the backups that are already stored and requires extra media.

Moving data for a client node

You can consolidate data for a client node by moving the data within server storage. You can move a backup set to different media, where the backup set is retained until the time that you specify. Consolidating data can help to improve efficiency during client restore or retrieve operations.

Data protection strategies with IBM Spectrum Protect

IBM Spectrum Protect™ provides ways for you to implement various data protection strategies.

You can configure IBM Spectrum Protect to send data to storage devices that are on the local site or on a remote site. To maximize data protection, you can configure replication to a remote server.

- Strategies to minimize the use of storage space for backups
To minimize the amount of storage space that is required, IBM Spectrum Protect backs up data by using the data deduplication and progressive incremental backup techniques.
- Strategies for disaster protection
IBM Spectrum Protect provides strategies to protect data if a disaster occurs. These strategies include node replication to a

remote site, storage pool protection, database backups, moving backup tapes offsite, and device replication to a standby server.

- Strategies for disaster recovery with IBM Spectrum Protect
IBM Spectrum Protect provides several ways to recover the server if the database or storage pools fail.


Strategies to minimize the use of storage space for backups

To minimize the amount of storage space that is required, IBM Spectrum Protect™ backs up data by using the data deduplication and progressive incremental backup techniques.

Data deduplication

When the IBM Spectrum Protect server receives data from a client, the server identifies duplicate data extents and stores unique instances of the data extents in a directory-container storage pool. The data deduplication technique improves storage utilization and eliminates the need for a dedicated data deduplication appliance.

Figure 1. Data deduplication process



If the same byte pattern occurs many times, data deduplication greatly reduces the amount of data that must be stored or transferred. In addition to whole files, IBM Spectrum Protect can also deduplicate parts of files that are common with parts of other files.

IBM Spectrum Protect provides the following types of data deduplication:

Server-side data deduplication

The server identifies duplicate data extents and moves the data to a directory-container storage pool. The server-side process uses *inline data deduplication*, where data is deduplicated at the same time that the data is written to a directory-container storage pool. Deduplicated data can also be stored in other types of storage pools. Inline data deduplication on the server provides the following benefits:

- Eliminates the need for reclamation
- Reduces the space that is occupied by the stored data

Client-side data deduplication

With this method, processing is distributed between the server and the client during a backup process. The client and the server identify and remove duplicate data to save storage space on the server. In client-side data deduplication, only compressed, deduplicated data is sent to the server. The server stores the data in the compressed format that is provided by the client. Client-side data deduplication provides the following benefits:

- Reduces the amount of data that is sent over the local area network (LAN)
- Eliminates extra processing power and time that is required to remove duplicate data on the server
- Improves database performance because the client-side data deduplication is also inline

You can combine both client-side and server-side data deduplication in the same production environment. The ability to deduplicate data on either the client or the server provides flexibility in terms of resource utilization, policy management, and data protection.

Compression

Use inline compression to reduce the amount of space that is stored in container storage pools. Data is compressed as it is written to the container storage pool.

Restriction: The IBM Spectrum Protect server cannot compress encrypted data.

Progressive incremental backup

In a progressive incremental backup process, the server monitors client activity and backs up any files that change since the initial full backup. Entire files are backed up, so that the server does not need to reference base versions of the files. This backup technique eliminates the need for multiple full backups of client data thus saving network resources and storage space.

Strategies for disaster protection

IBM Spectrum Protect™ provides strategies to protect data if a disaster occurs. These strategies include node replication to a remote site, storage pool protection, database backups, moving backup tapes offsite, and device replication to a standby server.

Replication to a remote site

Node replication is the process of incrementally copying data from one server to another server. The server from which client data is replicated is called a *source replication server*. The server to which client data is replicated is called a *target replication server*. For the purposes of disaster protection, the target replication server is on a remote site. A replication server can function as a source server, a target server, or both. You use replication processing to maintain the same level of files on the source and the target servers.

Node replication provides for immediate availability of data through failover. Although node replication protects most of the metadata, this approach does not provide adequate protection for database damage. You can provide more comprehensive protection by using storage pools to store data backups.

Advantages


- Failover so that data is available immediately if a disaster occurs.
- Incremental replication, which results in fast transmission of data.
- Electronic transfer
- Protects both data and most metadata

Disadvantages

- Both data and metadata must be recovered.
- Data on the source server must be replicated again from the remote site.

Figure 1 shows the node replication process to a remote site.

Figure 1. Node replication process



When client data is replicated, data that is not on the target server is copied to the target server. When replicated data exceeds the retention limit, the target server automatically removes the data from the source server. To maximize data protection, you synchronize the local server and the remote server; for example, Site B replicates data from Site A and Site A replicates data from Site B. As part of replication processing, client data that was deleted from the source server is also deleted from the target server.

IBM Spectrum Protect provides the following replication functions:

- You can define policies for the target server in the following ways:
 - Identical policies on the source server and target server
 - Different policies on the source server and target server to meet different business requirements.
 If a disaster occurs and the source server is not available, clients can recover data from the target server. If the source server cannot be recovered, you can direct clients to store data on the target server. When an outage occurs, the clients that are backed up to the source server can automatically fail over to restore their data from the target server.
- You can use replication processing to recover damaged files from storage pools. You must replicate the client data to the target server before the file damage occurs. Subsequent replication processes detect damaged files on the source server and replace the files with undamaged files from the target server.

Role of replication in disaster protection

If a disaster occurs, you can recover replicated data from the remote site and maintain the same level of files on the source and target servers. You use replication to achieve the following objectives:

- Control network throughput by scheduling node replication at specific times
- Recover data after a site loss.
- Recover damaged files on the source server.

Storage pool protection

As part of a disaster recovery strategy, ensure that a backup copy of data in storage pools is available at a remote site.

Advantages

- Fast recovery and rebuild of the source system.

Disadvantages

- Only data is protected; metadata is not protected.
- For each storage pool, you must define the storage medium.

You use different techniques to protect against the permanent loss of data that is stored in container storage pools and in FILE and DISK storage pools.

Directory-container storage pools

If you do not need to replicate all the data that is contained in a client node, you use container-copy storage pools to protect some directory-container storage pools. By protecting a directory-container storage pool, you do not use resources that replicate existing data and metadata, which improves server performance.

The preferred method is to protect the directory-container storage pool before you replicate the client node. When node replication is started, the data extents that are already replicated through storage pool protection are skipped, which reduces the replication processing time. If the data in a directory-container storage pool becomes damaged, you can repair the data from a copy in a container-copy storage pool.

Container-copy storage pools

You protect directory-container storage pools by copying the data in the directory-container storage pool to container-copy storage pools. Use container-copy storage pools to create up to two tape copies of a directory-container storage pool. The tape copies can be stored onsite or offsite. Damaged data in directory-container storage pools can be repaired by using container-copy storage pools. Container-copy storage pools provide an alternative to using a replication server to protect data in a directory-container storage pool.

Storage pools that are associated with FILE and DISK device classes

For storage pools that are associated with FILE and DISK device classes, you use node replication to maintain a node-consistent copy of the data at the target server. The data copy can be directly restored from the target server to the storage pools.

Database backups

You use database backups to recover your system following database damage. Also, database backup operations must be used to prevent DB2 from running out of archive log space. Database backup operations are not part of node replication. A database backup can be full, incremental, or snapshot. To provide for disaster recovery, a copy of the database backups must be stored offsite. To restore the database, you must have the backup volumes for the database. You can restore the database from backup volumes by either a point-in-time restore or a most current restore operation.

Point-in-time restore

Use point-in-time restore operations for situations such as disaster recovery or to remove the effects of errors that can cause inconsistencies in the database. Restore operations for the database that use snapshot backups are a form of point-in-time restore operation. The point-in-time restore operation includes the following actions:

- Removes and re-creates the active log directory and archive log directory that are specified in the dsmserv.opt file.
- Restores the database image from backup volumes to the database directories that are recorded in a database backup or to new directories.
- Restores archive logs from backup volumes to the overflow directory.
- Uses log information from the overflow directory up to a specified point in time.

Most current restore

If you want to recover the database to the time when the database was lost, recover the database to the most current state. The most current restore operation includes the following actions:

- Restores a database image from the backup volumes to the database directories that are recorded in a database backup or to new directories.
- Restores archive logs from backup volumes to the overflow directory.
- Uses log information from the overflow directory and archive logs from archive log directory.

The most current restore does not remove and re-create the active log directory or archive log directory.

Alternative methods for disaster protection

In addition to replication, storage pool protection, and database backups, you can also use the following methods to protect data and implement disaster recovery with IBM Spectrum Protect:

Sending backup tapes to a remote site

Data is backed up to tape at scheduled times by the source server. The tapes are sent to a remote site. If a disaster occurs, the tapes are returned to the site of the source server and the data is restored on the source clients. Offsite copies of data on backup tape can also help you to recover from ransomware attacks.

Multisite appliance replication to a standby server

In the multisite appliance configuration, the source appliance is replicated to a remote server in a SAN architecture. In this configuration, if the client hardware at the original site is damaged, the source device can be replicated from the standby server at the remote site. This configuration provides disk-based backup and restore operations.

Comparison of protection configuration strategies

Consider the following potential data-loss scenarios:

- Database data is damaged: protect against loss of data in the database by using onsite database backup.
- Storage pool data is damaged: protect against loss of data in storage pools by using onsite copy storage pools or node replication.
- Disaster scenario where both the onsite database and storage pools are lost: protect against a full disaster by using node replication and both off-site database backup and storage pool backup copies.

The following possible configurations address the most common data protection scenarios:

Configurations for damage protection only

- Implement database backup operations onsite with an optional container-copy storage pool onsite to protect data in directory-container storage pools.
- Implement database backup operations onsite and node replication onsite.

Configurations for disaster recovery and damage protection

- Implement database backup operations offsite with container-copy storage pools offsite to protect data in directory-container storage pools.
- Implement database backup operations onsite and node replication offsite with an optional container-copy storage pool onsite for faster recovery of damaged data.


Strategies for disaster recovery with IBM Spectrum Protect

IBM Spectrum Protect™ provides several ways to recover the server if the database or storage pools fail.

Automatic failover for disaster recovery

Automatic failover is an operation that switches to a standby system if a software, hardware, or network interruption occurs. Automatic failover is used with node replication to recover data after a system failure. Figure 1 shows the IBM Spectrum Protect automatic failover process.

Figure 1. Automatic failover process



Automatic failover for data recovery occurs if the source replication server is unavailable because of a disaster or a system outage. During normal operations, when the client accesses a source replication server, the client receives connection information for the target replication server. The client node stores the failover connection information in the client options file.

During client restore operations, the server automatically changes clients from the source replication server to the target replication server and back again. Only one server per node can be used for failover protection at any time. When a new client operation is started, the client attempts to connect to the source replication server. The client resumes operations on the source server if the source replication server is available.

To use automatic failover for replicated client nodes, the source replication server, the target replication server, and the client must be at the V7.1 level or later. If any of the servers are at an earlier level, automatic failover is disabled and you must rely on a manual failover process.

Recovery of IBM Spectrum Protect components

The server database, recovery log, and storage pools are critical to the operation of IBM Spectrum Protect and must be protected. If the database is unusable, the entire server is unavailable and recovering data that is managed by the server might be difficult or impossible.

Even without the database, fragments of data or complete files might be read from storage pool volumes that are not encrypted and security can be compromised. Therefore, you must always back up the database. Also, always encrypt sensitive data by using the client or the storage device, unless the storage media is physically secured.

IBM Spectrum Protect provides several data protection methods, which include backing up storage pools and the database. For example, you can define schedules so that the following operations occur:

- After the initial full backup of your storage pools, incremental storage pool backups are run every night.
- Incremental database backups are run every night.
- Full database backups are run once a week.

For tape-based environments, you can use disaster recovery manager (DRM) to assist you in many of the tasks that are associated with protecting and recovering data. DRM is available with IBM Spectrum Protect Extended Edition.

Preventive actions for recovery

Recovery is based on the following preventive actions:

- Mirroring, by which the server maintains a copy of the active log
- Backing up the database
- Backing up the storage pools

- Auditing storage pools for damaged files and recovery of damaged files when necessary
- Backing up the device configuration and volume history files
- Validating the data in storage pools by using cyclic redundancy checking
- Storing the cert.kdb file in a safe place to ensure that the Secure Sockets Layer (SSL) is secure

If you are using tape for storage, you can also create a disaster recovery plan to guide you through the recovery process by using DRM. You can use the disaster recovery plan for audit purposes to certify the recoverability of the server. The disaster recovery methods of DRM are based on taking the following actions:

- Creating a disaster recovery plan file for the server
- Backing up server data to tape
- Sending the server backup data to a remote site or to another server
- Storing client system information
- Defining and tracking the storage media that is used for storing and recovering client data

IBM Spectrum Protect data protection solutions

IBM Spectrum Protect™ servers and clients provide data protection solutions for the most common business and compliance requirements.

- Selecting a data protection solution for your environment
To help you to deploy a data protection environment, review information about best practice IBM Spectrum Protect configurations, and select the best solution for your business needs.
- Single-site disk solution
This data protection solution provides cost-effective data storage at a single site with minimal hardware setup.
- Multisite disk solution
This data protection solution provides replication at multiple sites so that each server protects data for the other site.
- Tape solution
This data protection solution provides storage to tape media, a flexible and affordable option for long-term data retention.
- Solutions documentation in PDF files
Prebuilt PDF files for IBM Spectrum Protect data protection solutions are available for you to download.

Selecting a data protection solution for your environment

To help you to deploy a data protection environment, review information about best practice IBM Spectrum Protect™ configurations, and select the best solution for your business needs.

- Disk-based implementation of a data protection solution for a single site
This disk-based implementation of a data protection solution with IBM Spectrum Protect uses inline data deduplication and provides protection for data on a single site.
- Disk-based implementation of a data protection solution for multiple sites
This disk-based implementation of a data protection solution with IBM Spectrum Protect uses inline data deduplication and replication at two sites.
- Appliance-based implementation of a data protection solution for multiple sites
This implementation of a multi-site IBM Spectrum Protect data protection solution uses appliance-based data deduplication and replication. A standby server is configured at a second site to recover data if the primary server is unavailable.
- Tape-based implementation of a data protection solution
This implementation of a data protection solution with IBM Spectrum Protect uses one or more tape storage devices to back up data. Tape backup provides low-cost scalability that is optimized for long-term retention.
- Comparison of data protection solutions
Compare the key features for each IBM Spectrum Protect solution to determine which configuration best meets your data protection requirements. Then, review the available documentation to implement the solution.
- Roadmap for implementing a data protection solution
Plan and implement the most suitable data protection solution for your business environment with IBM Spectrum Protect.

Disk-based implementation of a data protection solution for a single site

This disk-based implementation of a data protection solution with IBM Spectrum Protect™ uses inline data deduplication and provides protection for data on a single site.




Single-site disk

✓ Single-site architecture

✓ Cost effective

✓ Space efficient

✓ Simpler implementation



This data protection solution provides the following benefits:

- Server system and storage hardware at a single site
- Cost-effective use of storage through the data deduplication feature
- Space-efficient solution with minimal hardware setup
- Minimal implementation that requires installation and configuration for only one server and supporting storage hardware

In this solution, the client sends data to the IBM Spectrum Protect server, where the data is deduplicated and stored in a directory-container storage pool that is implemented in disk storage. Data from the inventory is also backed up to disk storage. This solution is suitable for entry-level environments for which a second copy of data is not required.


Related reference:

Comparison of data protection solutions

Roadmap for implementing a data protection solution

Disk-based implementation of a data protection solution for multiple sites

This disk-based implementation of a data protection solution with IBM Spectrum Protect™ uses inline data deduplication and replication at two sites.



This data protection solution provides the following benefits:

- Replication can be configured at both sites so that each server protects data for the other site
- Offsite data storage for each location is simplified
- Bandwidth is used efficiently because only deduplicated data is replicated between the sites
- Clients can automatically fail over to a target replication server if the source replication server is unavailable


In this solution, clients send data to the source server, where the data is deduplicated and stored in a directory-container storage pool that is implemented in disk storage. The data is replicated to the storage pool on the target server for each site. This solution is suitable for environments that require disaster protection. If mutual replication is configured, clients at both sites can use failover recovery for continued backups and data recovery from the available server on the other site.

Related reference:

- Comparison of data protection solutions
- Roadmap for implementing a data protection solution

Appliance-based implementation of a data protection solution for multiple sites

This implementation of a multi-site IBM Spectrum Protect™ data protection solution uses appliance-based data deduplication and replication. A standby server is configured at a second site to recover data if the primary server is unavailable.



This data protection solution provides the following benefits:

- Performance is optimized for backups on high-speed storage area networks (SAN) and for use with IBM Spectrum Protect for SAN, when clients back up directly to SAN-attached virtual tape devices.
- Fast, appliance-based replication frees the server from having to track replication metadata in the server database.
- Bandwidth and storage space are used efficiently because only deduplicated data is replicated between the sites.
- A standby environment provides for disaster recovery, but does not require the amount of resources that are needed for a fully active site.

In this data protection configuration, the server uses hardware appliances to deduplicate and replicate data. The appliance at Site A deduplicates data and then replicates the data to the appliance at Site B for disaster protection. If a failure at Site A occurs, you make the standby server active by restoring the most recent database backup, and by activating the replicated copy of data.


For more information about configuring virtual tape libraries, see [Configuring virtual tape libraries](#).

Related reference:

- Comparison of data protection solutions
- Roadmap for implementing a data protection solution

Tape-based implementation of a data protection solution

This implementation of a data protection solution with IBM Spectrum Protect™ uses one or more tape storage devices to back up data. Tape backup provides low-cost scalability that is optimized for long-term retention.



This data protection solution provides the following benefits:

- Performance is optimized for backup operations on high-speed storage area networks (SAN) directly to tape for large data types and for long-term retention of data.
- Data availability is optimized by storing copies of data at offsite locations for disaster recovery. If you enable the disaster recovery management (DRM) function and a disaster occurs, DRM helps to streamline the process of recovering your servers.
- Data security is optimized because copies of data are stored offsite on tape devices that are *not* connected to the internet. Ransomware attacks rely on internet connections; therefore, offsite storage can help to protect against such attacks.
- Low-cost scalability is achieved by reducing the need for additional disk hardware and lowering energy costs.

Related concepts:

Selecting a tape device driver

Related tasks:

Creating data backup strategies

Managing volume inventory

Related reference:

Comparison of data protection solutions

Installing and configuring tape device drivers

Comparison of data protection solutions

Compare the key features for each IBM Spectrum Protect™ solution to determine which configuration best meets your data protection requirements. Then, review the available documentation to implement the solution.

	Single-site disk	Multisite disk	Multisite appliance	Tape
Highlights				
Cost	\$	\$\$\$	\$\$\$\$	\$\$
Protection level	One data copy	Two or more data copies	Two or more data copies	Two or more data copies

	Single-site disk	Multisite disk	Multisite appliance	Tape
				
Disaster recovery	None	Active server	Standby server	Offsite copies
Key benefits				
Leading-edge data reduction	✓	✓	✓	✓
Fast and efficient disk-based backup and restore operations	✓		✓	
Simplified offsite management		✓		
Data deduplication feature at no extra cost	✓	✓		
Replication processing included at no extra charge		✓		
Data deduplication at both the source and target server		✓		
Low-cost scalability and optimized for long-term retention				✓
Efficiency and cost				
Optimized for high-speed storage area network (SAN) backup operations			✓	✓
Optimized for high-speed local area network (LAN)	✓	✓	✓	
Global data deduplication across all data types and sources	✓	✓	✓	
Bandwidth-efficient replication		✓	✓	
Lower energy costs				✓
Option for a second copy without extra disk hardware				✓
Availability				
Offsite copy capability		✓	✓	✓
Appliance-based replication			✓	
Client recovery from high-availability server		✓		
Replication target in the cloud		✓		
Independent management of retention policies for replication data; ability to keep more or less data at recovery site		✓		
Application-level replication; ability to choose which systems and applications are replicated		✓		
Scalability				
Global data deduplication across servers			✓	
SAN-optimized backup directly to tape for large data types				✓
Single-instance petabyte scalability				✓

What to do next

Review available documentation for the solutions in Roadmap for implementing a data protection solution.

Related reference:

Disk-based implementation of a data protection solution for a single site

Disk-based implementation of a data protection solution for multiple sites

Appliance-based implementation of a data protection solution for multiple sites

Tape-based implementation of a data protection solution

Roadmap for implementing a data protection solution

Plan and implement the most suitable data protection solution for your business environment with IBM Spectrum Protect™.

Single-site disk solution

For steps that describe how to plan for, implement, monitor, and operate a single-site disk solution, see [Single-site disk solution](#).

Multisite disk solution

For steps that describe how to plan for, implement, monitor, and operate a multisite disk solution, see [Multisite disk solution](#).

Tape solution

For steps that describe how to plan for, implement, monitor, and operate a tape device solution, see [Tape solution](#).

Multisite appliance solution

For an overview of the tasks that are required to implement a multisite appliance solution, review the following steps:

1. Begin planning for the solution by reviewing information at the following links:
 - o AIX: Capacity planning
 - o Linux: Capacity planning
 - o Windows: Capacity planning
2. Install the server and optionally, the Operations Center. Review information at the following links:
 - o Installing and upgrading the server
 - o Installing and upgrading the Operations Center
3. Configure the server for storage in a virtual tape library.
 - o Managing virtual tape libraries
 - o Attaching tape devices for the server

For guidance about improving system performance, see [Configuration best practices](#).

4. Configure policies to protect your data. Review the information in [Customizing policies](#).
5. Set up client schedules. Review the information in [Scheduling backup and archive operations](#).
6. Install and configure clients. To determine the type of client software that you need, review the information in [Adding clients for details](#).
7. Configure monitoring for your system. Review the information in [Monitoring storage solutions](#).

Related reference:

Comparison of data protection solutions

Disk-based implementation of a data protection solution for a single site

Disk-based implementation of a data protection solution for multiple sites

Appliance-based implementation of a data protection solution for multiple sites

Tape-based implementation of a data protection solution

Single-site disk solution

This data protection solution provides cost-effective data storage at a single site with minimal hardware setup.

- Planning for a single-site disk data protection solution
 - Plan for a data protection implementation that includes a server at a single site that uses data deduplication.

- Single-site disk implementation of a data protection solution
The single-site disk solution is configured at one site and uses data deduplication.
- Monitoring a single-site disk solution
After you implement a single-site disk solution with IBM Spectrum Protect, monitor the solution for correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.
- Managing operations for a single-site disk solution
Use this information to manage operations for a single-site disk solution with IBM Spectrum Protect that includes a server and uses data deduplication for a single location.

Planning for a single-site disk data protection solution

Plan for a data protection implementation that includes a server at a single site that uses data deduplication.

Implementation options

You can configure the server for a single-site disk solution in the following ways:

Configure the server by using the Operations Center and administrative commands

This documentation provides steps to configure a range of storage systems and the server software for your solution. Configuration tasks are completed by using wizards and options in the Operations Center and IBM Spectrum Protect™ commands. For information about getting started, see the Planning roadmap.

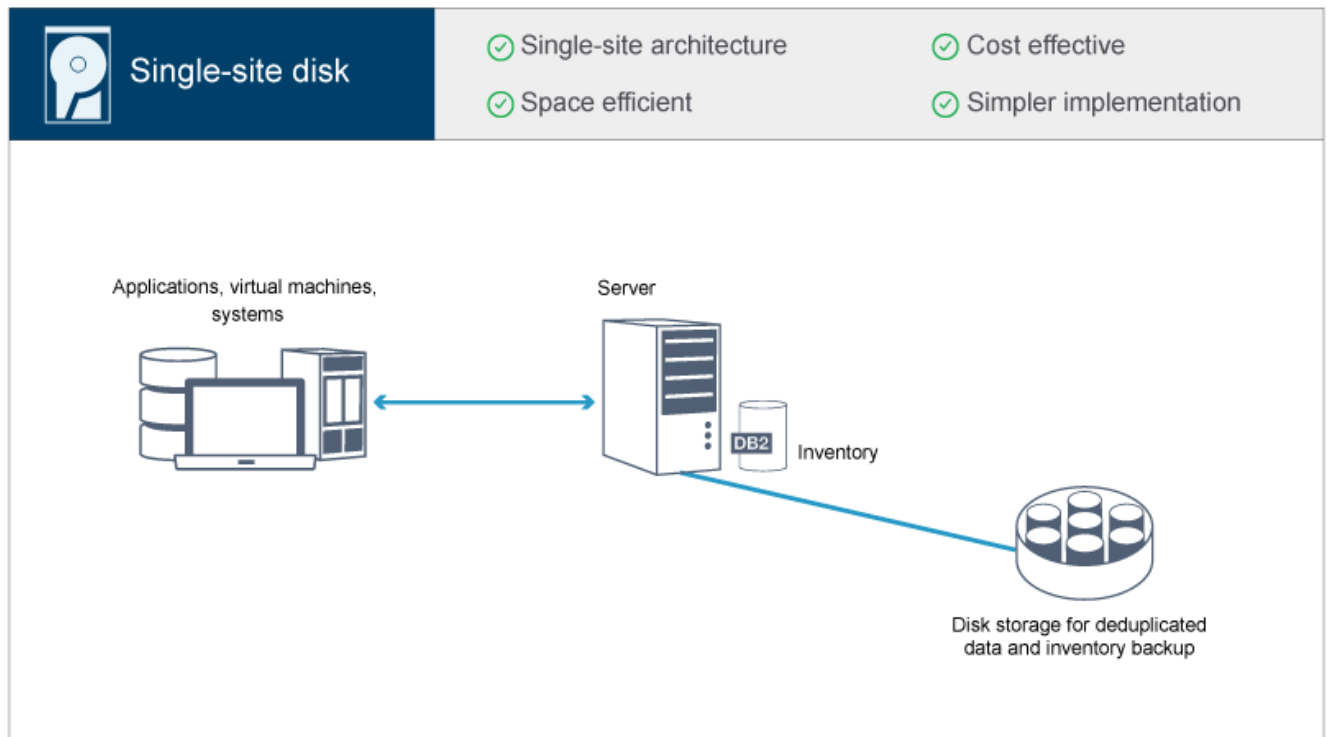
Configure the server by using automated scripts

For detailed guidance on implementing a single-site disk solution with specific IBM® Storwize® storage systems, and by using automated scripts to configure the server, see the IBM Spectrum Protect blueprints. The documentation and scripts are available on IBM developerWorks® at: IBM Spectrum Protect Blueprints.

The blueprint documentation does not include steps for installing and configuring the Operations Center, or setting up secure communications by using Transport Security Layer (TLS). An option for using Elastic Storage Server, based on IBM Spectrum Scale™ technology, is included.

Planning roadmap

Plan for the single-site disk solution by reviewing the architecture layout in the following figure and then completing the roadmap tasks that follow the diagram.



The following steps are required to plan for a single-site disk environment.

1. Select your system size.
2. Meet system requirements for hardware and software.
3. Record values for your system configuration in the planning worksheets.
4. Plan for storage.
5. Plan for security.
 - a. Plan for administrator roles.
 - b. Plan for secure communications.
 - c. Plan for storage of encrypted data.
 - d. Plan for firewall access.

Selecting a system size

Select the size of the IBM Spectrum Protect™ server based on the amount of data that you manage and the systems to be protected.

About this task

You can use the information in the table to determine the size of the server that is required, based on the amount of data that you manage.

The following table describes the volume of data that a server manages. This amount includes all versions. The daily amount of data is how much new data you back up each day. Both the total managed data and daily amount of new data are measured as the size before any data reduction.

Table 1. Determining the size of the server

Total managed data	Daily amount of new data to back up	Required server size
60 TB - 240 TB	Up to 10 TB per day	Small
196 TB - 784 TB	10 - 20 TB per day	Medium
1000 TB - 4000 TB	20 - 100 TB per day	Large

The daily backup values in the table are based on test results with 128 MB sized objects, which are used by IBM Spectrum Protect for Virtual Environments. Workloads that consist of objects that are smaller than 128 KB might not be able to achieve these daily limits.

System requirements for a single-site disk solution

After you select the IBM Spectrum Protect™ solution that best fits your data protection requirements, review the system requirements to plan for implementation of the data protection solution.

Ensure that your system meets the hardware and software prerequisites for the size of server that you plan to use.

- **Hardware requirements**
Hardware requirements for your IBM Spectrum Protect solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.
- **Software requirements**
Documentation for the single-site disk IBM Spectrum Protect solution includes installation and configuration tasks for the following operating systems. You must meet the minimum software requirements that are listed.

Related information:

[IBM Spectrum Protect Supported Operating Systems](#)

Hardware requirements

Hardware requirements for your IBM Spectrum Protect™ solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.

For a definition of system sizes, see [t_ssdisk_select_size.html](#).

The following table includes minimum hardware requirements for the server and storage, based on the size of the server that you plan to build. If you are using local partitions (LPARs) or work partitions (WPARs), adjust the network requirements to take

account of the partition sizes.

Use the information in the following table as a starting point. For the most up-to-date information about hardware requirements and specifications for the server and storage, see the IBM Spectrum Protect Blueprints.

Hardware component	Small system	Medium system	Large system
Server processor	<p>AIX 6 processor cores, 3.42 GHz or faster</p> <p>Linux Windows 16 processor cores, 1.7 GHz or faster</p>	<p>AIX 10 processor cores, 3.42 GHz or faster</p> <p>Linux Windows 20 processor cores, 2.2 GHz or faster</p>	<p>AIX 20 processor cores, 3.42 GHz</p> <p>Linux Windows 44 processor cores, 2.2 GHz or faster</p>
Server memory	64 GB RAM	128 GB RAM	256 GB RAM
Network	<ul style="list-style-type: none"> 10 GB Ethernet (1 port) 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> 10 GB Ethernet (2 ports) 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> 10 GB Ethernet (4 ports) 8 GB Fibre Channel adapter (4 ports)
Storage	<ul style="list-style-type: none"> 1.45 TB SSD disks for the database, plus space for Operations Center records 67 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> 2.53 TB SSD disks for the database, plus space for Operations Center records 207.9 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> 6.54 TB SSD disks for the database, plus space for Operations Center records 1049.67 TB deduplicated directory-container storage pool

Estimating database space requirements for the Operations Center

Hardware requirements for the Operations Center are included in the preceding table, except for the database and archive log space (inventory) that the Operations Center uses to hold records for managed clients.

If you do not plan to install the Operations Center on the same system as the server, you can estimate system requirements separately. To calculate system requirements for the Operations Center, see the system requirements calculator in technote 1641684.

Managing the Operations Center on the server is a workload that requires extra space for database operations. The amount of space depends on the number of clients that are monitored on a server. Review the following guidelines to estimate how much space your server requires.

Database space

The Operations Center uses approximately 1.2 GB of database space for every 1000 clients that are monitored on a server. For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1500 clients. This configuration has a total of 6500 clients across the four servers and requires approximately 8.4 GB of database space. This value is calculated by rounding the 6500 clients up to the next closest 1000, which is 7000:

$$7 \times 1.2 \text{ GB} = 8.4 \text{ GB}$$

Archive log space

The Operations Center uses approximately 8 GB of archive log space every 24 hours, for every 1000 clients. In the example of 6500 clients across the hub server and the spoke servers, 56 GB of archive log space is used over a 24-hour period for the hub server.

For each spoke server in the example, the archive log space that is used over 24 hours is approximately 16 GB. These estimates are based on the default status collection interval of 5 minutes. If you reduce the collection interval from once every 5 minutes to once every 3 minutes, the space requirements increase. The following examples show the approximate increase in the log space requirement with a collection interval of once every 3 minutes:

- Hub server: 56 GB to approximately 94 GB
- Each spoke server: 16 GB to approximately 28 GB

Increase the archive log space so that you have sufficient space available to support the Operations Center, without affecting the existing server operations.

Software requirements

Documentation for the single-site disk IBM Spectrum Protect™ solution includes installation and configuration tasks for the following operating systems. You must meet the minimum software requirements that are listed.

For information about software requirements for IBM® lin_tape device drivers, refer to the IBM Tape Device Drivers Installation and User's Guide.

AIX systems

Type of software	Minimum software requirements
Operating system	IBM AIX® 7.1 For more information about operating system requirements, see AIX: Minimum system requirements for AIX systems.
Gunzip utility	The gunzip utility must be available on your system before you install or upgrade the IBM Spectrum Protect server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.
File system type	JFS2 file systems AIX systems can cache a large amount of file system data, which can reduce memory that is required for server and IBM DB2® processes. To avoid paging with the AIX server, use the rbrw mount option for the JFS2 file system. Less memory is used for the file system cache and more is available for IBM Spectrum Protect. Do not use the file system mount options, Concurrent I/O (CIO), and Direct I/O (DIO), for file systems that contain the IBM Spectrum Protect database, logs, or storage pool volumes. These options can cause performance degradation of many server operations. IBM Spectrum Protect and DB2 can still use DIO where it is beneficial to do so, but IBM Spectrum Protect does not require the mount options to selectively take advantage of these techniques.
Other software	Korn Shell (ksh)

Linux systems

Type of software	Minimum software requirements
Operating system	Red Hat Enterprise Linux 7 (x86_64)
Libraries	GNU C libraries, Version 2.3.3-98.38 or later that is installed on the IBM Spectrum Protect system. Red Hat Enterprise Linux Servers: <ul style="list-style-type: none">• libaio• libstdc++.so.6 (32-bit and 64-bit packages are required)• numactl.x86_64
File system type	Format database-related file systems with ext3 or ext4. For storage pool-related file systems, use XFS.
Other software	Korn Shell (ksh)

Windows systems

Type of software	Minimum software requirements
Operating system	Microsoft Windows Server 2012 R2 (64-bit) or Windows Server 2016
File system type	NTFS

Type of software	Minimum software requirements
Other software	<p>Windows 2012 R2 or Windows 2016 with .NET Framework 3.5 is installed and enabled.</p> <p>The following User Account Control policies must be disabled:</p> <ul style="list-style-type: none"> • User Account Control: Admin Approval Mode for the Built-in Administrator account • User Account Control: Run all administrators in Admin Approval Mode

Related tasks:

➔ [Setting AIX network options](#)

Planning worksheets

Use the planning worksheets to record values that you use to set up your system and configure the IBM Spectrum Protect™ server. Use the best practice default values that are listed in the worksheets.

Each worksheet helps you prepare for different parts of the system configuration by using best practice values:

Server system preconfiguration

Use the preconfiguration worksheets to plan for the file systems and directories that you create when you configure file systems for IBM Spectrum Protect during system setup. All directories that you create for the server must be empty.

Server configuration

Use the configuration worksheets when you configure the server. Default values are suggested for most items, except where noted.



Table 1. Worksheet for preconfiguration of an AIX server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	/home/tsminst1/tsminst1		50 GB	If you change the value for the server instance directory from the default, also modify the DB2® instance owner value in Table 2.
Directory for server installation	/		Available space that is required for the directory: 5 GB	
Directory for server installation	/usr		Available space that is required for the directory: 5 GB	
Directory for server installation	/var		Available space that is required for the directory: 5 GB	
Directory for server installation	/tmp		Available space that is required for the directory: 5 GB	

Item	Default value	Your value	Minimum directory size	Notes
Directory for server installation	/opt		Available space that is required for the directory: 10 GB	
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	When you create the active log during the initial configuration of the server, set the size to 128 GB.
Directory for the archive log	/tsminst1/TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 ...		<p>Minimum total space for all directories:</p> <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	<p>Create a minimum number of file systems for the database, depending on the size of your system:</p> <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		<p>Minimum total space for all directories:</p> <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	<p>Create a minimum number of file systems for storage, depending on the size of your system:</p> <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 <p>The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.</p>

Table 2. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 1 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	passw0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Item	Default value	Your value	Notes
Schedule start time	22:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Linux

Table 3. Worksheet for preconfiguration of a Linux server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	<p>Ensure that this port is available when you install and configure the operating system</p> <p>The port number can be a number in the range 1024 - 32767.</p>
Directory for the server instance	/home/tsminst1/tsminst1		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 4.
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	
Directory for the archive log	/tsminst1/TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	

Item	Default value	Your value	Minimum directory size	Notes
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	Create a minimum number of file systems for the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 <p>The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.</p>

Table 4. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
------	---------------	------------	-------

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 3 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	passw0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.
Schedule start time	22:00		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window. Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Windows

Because many volumes are created for the server, configure the server by using the Windows feature of mapping disk volumes to directories rather than to drive letters.

For example, C:\tsminst1\TSMdbpspace00 is a mount point to a volume with its own space. The volume is mapped to a directory under the C: drive, but does not take up space from the C: drive. The exception is the server instance directory, C:\tsminst1, which can be a mount point or a regular directory.

Table 5. Worksheet for preconfiguration of a Windows server system

Item	Default value	Your value	Minimum directory size	Notes
------	---------------	------------	------------------------	-------

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	C:\tsminst1		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 6.
Directory for the active log	C:\tsminst1\TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	
Directory for the archive log	C:\tsminst1\TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	C:\tsminst1\TSMdbspace00 C:\tsminst1\TSMdbspace01 C:\tsminst1\TSMdbspace02 C:\tsminst1\TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	Create a minimum number of file systems for the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	C:\tsminst1\TSMbkup00 C:\tsminst1\TSMbkup01 C:\tsminst1\TSMbkup02 C:\tsminst1\TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Table 6. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 5 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	pAssW0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Item	Default value	Your value	Notes
Schedule start time	22:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Planning for storage

Choose the most effective storage technology for IBM Spectrum Protect™ components to ensure efficient server performance and operations.

Storage hardware devices have different capacity and performance characteristics, which determine how they can be used effectively with IBM Spectrum Protect. For general guidance on selecting the appropriate storage hardware and set up for your solution, review the following guidelines.

Database and active log

- Use a fast disk for the IBM Spectrum Protect database and active log, for example with the following characteristics:
 - High performance, 15k rpm disk with Fibre Channel or serial-attached SCSI (SAS) interface
 - Solid-state disk (SSD)
- Isolate the active log from the database unless you use SSD or flash hardware
- When you create arrays for the database, use RAID level 5

Storage pool

- You can use less expensive and slower disks for the storage pool
- The storage pool can share disks for the archive log and database backup storage
- Use RAID level 6 for storage pool arrays to add protection against double drive failures when you use large disk types
- Planning the storage arrays

Prepare for disk storage configuration by planning for RAID arrays and volumes, according to the size of your IBM Spectrum Protect system.

Related reference:

[Storage system requirements and reducing the risk of data corruption](#)

Planning for security

Plan to protect the security of systems in the IBM Spectrum Protect™ solution with access and authentication controls, and consider encrypting data and password transmission.

For guidelines about protecting your storage environment against ransomware attacks, and recovering your storage environment if an attack occurs, see Protecting the storage environment against ransomware.

- Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect solution.
- Planning for secure communications

Plan for protecting communications among the IBM Spectrum Protect solution components.

- Planning for storage of encrypted data
Determine whether your company requires stored data to be encrypted, and choose the option that best suits your needs.
- Planning firewall access
Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect solution to work.

Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect™ solution.

You can assign one of the following levels of authority to administrators:

System

Administrators with system authority have the highest level of authority. Administrators with this level of authority can complete any task. They can manage all policy domains and storage pools, and grant authority to other administrators.

Policy

Administrators who have policy authority can manage all of the tasks that are related to policy management. This privilege can be unrestricted, or can be restricted to specific policy domains.

Storage

Administrators who have storage authority can allocate and control storage resources for the server.

Operator

Administrators who have operator authority can control the immediate operation of the server and the availability of storage media such as tape libraries and drives.

The scenarios in Table 1 provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

Table 1. Scenarios for administrator roles

Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	<ul style="list-style-type: none"> • System authority: 1 administrator ID
An administrator for multiple servers also manages the overall system. Several other administrators manage their own storage pools.	<ul style="list-style-type: none"> • System authority on all servers: 1 administrator ID for the overall system administrator • Storage authority for designated storage pools: 1 administrator ID for each of the other administrators
An administrator manages 2 servers. Another person helps with the administration tasks. Two assistants are responsible for helping to ensure that important systems are backed up. Each assistant is responsible for monitoring the scheduled backups on one of the IBM Spectrum Protect servers.	<ul style="list-style-type: none"> • System authority on both servers: 2 administrator IDs • Operator authority: 2 administrator IDs for the assistants with access to the server that each person is responsible for

Planning for secure communications

Plan for protecting communications among the IBM Spectrum Protect™ solution components.

Determine the level of protection that is required for your data, based on regulations and business requirements under which your company operates.

If your business requires a high level of security for passwords and data transmission, plan on implementing secure communication with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocols.

TLS and SSL provide secure communications between the server and client, but can affect system performance. To improve system performance, use TLS for authentication without encrypting object data. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option for client-to-server communication, and the UPDATE SERVER=SSL parameter for server-to-server communication. Beginning in V8.1.2, TLS is used for authentication by default. If you decide to use TLS to encrypt entire sessions, use the protocol only for sessions where it is necessary and add processor resources on the server to manage the increase in network traffic. You can also try other options. For example, some networking devices such as routers and switches provide the TLS or SSL function.

You can use TLS and SSL to protect some or all of the different possible communication paths, for example:

- Operations Center: browser to hub; hub to spoke
- Client to server
- Server to server: node replication

Related tasks:

[🔗 Securing communications](#)

Planning for storage of encrypted data

Determine whether your company requires stored data to be encrypted, and choose the option that best suits your needs.

If your company requires the data in storage pools to be encrypted, then you have the option of using IBM Spectrum Protect™ encryption, or an external device such as tape for encryption.

If you choose IBM Spectrum Protect to encrypt the data, extra computing resources are required at the client that might affect the performance of backup and restore processes.

Related information:

[🔗 technote 1963635](#)

Planning firewall access

Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect™ solution to work.

Table 1 describes the ports that are used by the server, client, and Operations Center.

Table 1. Ports that are used by the server, client, and Operations Center

Item	Default	Direction	Description
Base port (TCPSPORT)	1500	Outbound/inbound	Each server instance requires a unique port. You can specify an alternative port number instead of using the default. The TCPSPORT option listens for both TCP/IP and SSL-enabled sessions from the client. For administrative client traffic, you can use the TCPADMINPORT and ADMINONCLIENTPORT options to set port values.
SSL-only port (SSLTCPSPORT)	No default	Outbound/inbound	This port is used if you want to restrict communication on the port to SSL-enabled sessions only. To support both SSL and non-SSL communications, use the TCPSPORT or TCPADMINPORT options.
SMB	45	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SSH	22	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SMTP	25	Outbound	This port is used to send email alerts from the server.
NDMP	No default	Inbound/outbound	<p>The server must be able to open an outbound NDMP control port connection to the NAS device. The outbound control port is the Low-Level Address in the data mover definition for the NAS device.</p> <p>During an NDMP filer-to-server restore, the server must be able to open an outbound NDMP data connection to the NAS device. The data connection port that is used during a restore can be configured on the NAS device.</p> <p>During NDMP filer-to-server backups, the NAS device must be able to open outbound data connections to the server and the server must be able to accept inbound NDMP data connections. You can use the server option NDMPPORTRANGE to restrict the set of ports available for use as NDMP data connections. You can configure a firewall for connections to these ports.</p>

Item	Default	Direction	Description
Replication	No default	Outbound/inbound	The port and protocol for the outbound port for replication are set by the DEFINE SERVER command that is used to set up replication. The inbound ports for replication are the TCP ports and SSL ports that the source server names in the DEFINE SERVER command.
Client schedule port	Client port: 1501	Outbound	The client listens on the port that is named and communicates the port number to the server. The server contacts the client if server prompted scheduling is used. You can specify an alternative port number in the client options file.
Long running sessions	KEEPALIVE setting: YES	Outbound	When the KEEPALIVE option is enabled, keepalive packets are sent during client-server sessions to prevent the firewall software from closing long-running, inactive connections.
Operations Center	HTTPS: 11090	Inbound	These ports are used for the Operations Center web browser. You can specify an alternative port number.
Client management service port	Client port: 9028	Inbound	The client management service port must be accessible from the Operations Center. Ensure that firewalls cannot prevent connections. The client management service uses the TCP port of the server for the client node for authentication by using an administrative session.

Single-site disk implementation of a data protection solution

The single-site disk solution is configured at one site and uses data deduplication.

Implementation roadmap

The following steps are required to set up the IBM Spectrum Protect™ single-site disk environment.

1. Set up the system.
 - a. Configure the storage hardware and set up storage arrays for your environment size.
 - b. Install the server operating system.
 - c. Configure multipath I/O.
 - d. Create the user ID for the server instance.
 - e. Prepare file systems for IBM Spectrum Protect.
2. Install the server and Operations Center.
3. Configure the server and Operations Center.
 - a. Complete the initial configuration of the server.
 - b. Set server options.
 - c. Configure Secure Sockets Layer for the server and client.
 - d. Configure the Operations Center.
 - e. Register your IBM Spectrum Protect license.
 - f. Configure data deduplication.
 - g. Define data retention rules for your business.
 - h. Define server maintenance schedules.
 - i. Define client schedules.
4. Install and configure clients.
 - a. Register and assign clients to schedules.
 - b. Install and verify the client management service.
 - c. Configure the Operations Center to use the client management service.
5. Complete the implementation.

Setting up the system

To set up the system, you must first configure your disk storage hardware and the server system for IBM Spectrum Protect™.

- Configuring the storage hardware
To configure your storage hardware, review general guidance for disk systems and IBM Spectrum Protect.
- Installing the server operating system
Install the operating system on the server system and ensure that IBM Spectrum Protect server requirements are met. Adjust operating system settings as directed.
- Configuring multipath I/O
You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.
- Creating the user ID for the server
Create the user ID that owns the IBM Spectrum Protect server instance. You specify this user ID when you create the server instance during initial configuration of the server.
- Preparing file systems for the server
You must complete file system configuration for the disk storage to be used by the server.

Configuring the storage hardware

To configure your storage hardware, review general guidance for disk systems and IBM Spectrum Protect™.

Procedure

1. Provide a connection between the server and the storage devices by following these guidelines:
 - Use a switch or direct connection for Fibre Channel connections.
 - Consider the number of ports that are connected and account for the amount of bandwidth that is needed.
 - Consider the number of ports on the server and the number of host ports on the disk system that are connected.
2. Verify that device drivers and firmware for the server system, adapters, and operating system are current and at the recommended levels.
3. Configure storage arrays. Make sure that you planned properly to ensure optimal performance. See Planning for storage for more information.
4. Ensure that the server system has access to disk volumes that are created. Complete the following steps:
 - a. If the system is connected to a Fibre Channel switch, zone the server to see the disks.
 - b. Map all of the volumes to tell the disk system that this specific server is allowed to see each disk.

Installing the server operating system

Install the operating system on the server system and ensure that IBM Spectrum Protect™ server requirements are met. Adjust operating system settings as directed.

- Installing on AIX systems
Complete the following steps to install AIX® on the server system.
- Installing on Linux systems
Complete the following steps to install Linux x86_64 on the server system.
- Installing on Windows systems
Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect server.

Installing on AIX systems

Complete the following steps to install AIX® on the server system.

Procedure

1. Install AIX Version 7.1, TL4, SP2, or later according to the manufacturer instructions.
2. Configure your TCP/IP settings according to the operating system installation instructions.
3. Open the /etc/hosts file and complete the following actions:
 - Update the file to include the IP address and host name for the server. For example:


```
192.0.2.7 server.yourdomain.com server
```
 - Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:


```
127.0.0.1 localhost
```

4. Enable AIX I/O completion ports by issuing the following command:

```
chdev -l iocp0 -P
```

Server performance can be affected by the Olson time zone definition.

5. To optimize performance, change your system time zone format from Olson to POSIX. Use the following command format to update the time zone setting:

```
chtz=local_timezone,date/time,date/time
```

For example, if you lived in Tucson, Arizona, where Mountain Standard Time is used, you would issue the following command to change to the POSIX format:

```
chtz MST7MDT,M3.2.0/2:00:00,M11.1.0/2:00:00
```

6. Add an entry in the .profile of the instance user so that the following environment is set:

```
export MALLOCOPTIONS=multiheap:16
```

Tip: If the instance user is not available, complete this step later, when the instance user becomes available.

7. Set the system to create full application core files. Issue the following command:

```
chdev -l sys0 -a fullcore=true -P
```

8. For communications with the server and Operations Center, make sure that the following ports are open on any firewalls that might exist:

- o For communications with the server, open port 1500.
- o For secure communications with the Operations Center, open port 11090 on the hub server.

If you are not using the default port values, make sure that the ports that you are using are open.

9. Enable TCP high-performance enhancements. Issue the following command:

```
no -p -o rfc1323=1
```

10. For optimal throughput and reliability, bond four 10 Gb Ethernet ports together. Use the System Management Interface Tool (SMIT) to bond the ports together by using Etherchannel. The following settings were used during testing:

```

mode          8023ad          Enable automatic recovery after failover
auto_recovery yes            Adapter used when whole channel fails
backup_adapter NONE          Determines how outgoing adapter is chosen
hash_mode     src_dst_port   Determines interval value for IEEE
interval      long           802.3ad mode

mode          8023ad          EtherChannel mode of operation
netaddr       0              Address to ping
no_loss_failover yes        Enable lossless failover after ping
failure

num_retries   3              Times to retry ping before failing
retry_time    1              Wait time (in seconds) between pings
use_alt_addr  no             Enable Alternate EtherChannel Address
use_jumbo_frame no          Enable Gigabit Ethernet Jumbo Frames

```

11. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 1. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 1. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf
Maximum number of open files	nofile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht

User limit type	Setting	Value	Command to query value
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Linux systems

Complete the following steps to install Linux x86_64 on the server system.

Before you begin

The operating system will be installed on the internal hard disks. Configure the internal hard disks by using a hardware RAID 1 array. For example, if you are configuring a small system, the two 300 GB internal disks are mirrored in RAID 1 so that a single 300 GB disk appears available to the operating system installer.

Procedure

1. Install Red Hat Enterprise Linux Version 7.1 or later, according to the manufacturer instructions. Obtain a bootable DVD that contains Red Hat Enterprise Linux Version 7.1 and start your system from this DVD. See the following guidance for installation options. If an item is not mentioned in the following list, leave the default selection.
 - a. After you start the DVD, choose Install or upgrade an existing system from the menu.
 - b. On the Welcome screen, select Test this media & install Red Hat Enterprise Linux 7.1.
 - c. Select your language and keyboard preferences.
 - d. Select your location to set the correct time zone.
 - e. Select Software Selection and then on the next screen, select Server with GUI.
 - f. From the installation summary page, click Installation Destination and verify the following items:
 - The local 300 GB disk is selected as the installation target.
 - Under Other Storage Options, Automatically configure partitioning is selected.
 - Click Done.
 - g. Click Begin Installation. After the installation starts, set the root password for your root user account.

After the installation is completed, restart the system and log in as the root user. Issue the `df` command to verify your basic partitioning. For example, on a test system, the initial partitioning produced the following result:

```
[root@tvapp02]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  50G  3.0G  48G   6% /
devtmpfs        32G   0    32G   0% /dev
tmpfs           32G  92K   32G   1% /dev/shm
tmpfs           32G  8.8M  32G   1% /run
tmpfs           32G   0    32G   0% /sys/fs/cgroup
/dev/mapper/rhel-home 220G  37M  220G   1% /home
/dev/sda1       497M 124M  373M  25% /boot
```

2. Configure your TCP/IP settings according to the operating system installation instructions.

For optimal throughput and reliability, consider bonding multiple network ports together. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection, which aggregates several subordinate ports into a single logical connection. The preferred method is to use a bond mode of 802.3ad, miimon setting of 100, and a `xmit_hash_policy` setting of layer3+4.

Restriction: To use an LACP network connection, you must have a network switch that supports LACP.

For additional instructions about configuring bonded network connections with Red Hat Enterprise Linux Version 7, see [Create a Channel Bonding Interface](#).
3. Open the `/etc/hosts` file and complete the following actions:
 - Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7  server.yourdomain.com  server
```
 - Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1  localhost
```
4. Install components that are required for the server installation. Complete the following steps to create a Yellowdog Updater Modified (YUM) repository and install the prerequisite packages.

- a. Mount your Red Hat Enterprise Linux installation DVD to a system directory. For example, to mount it to the /mnt directory, issue the following command:

```
mount -t iso9660 -o ro /dev/cdrom /mnt
```

- b. Verify that the DVD mounted by issuing the mount command. You should see output similar to the following example:

```
/dev/sr0 on /mnt type iso9660
```

- c. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

If the repos.d directory does not exist, create it.

- d. List directory contents:

```
ls rhel-source.repo
```

- e. Rename the original repo file by issuing the mv command. For example:

```
mv rhel-source.repo rhel-source.repo.orig
```

- f. Create a new repo file by using a text editor. For example, to use the vi editor, issue the following command:

```
vi rhel71_dvd.repo
```

- g. Add the following lines to the new repo file. The baseurl parameter specifies your directory mount point:

```
[rhel71_dvd]
name=DVD Redhat Enterprise Linux 7.1
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

- h. Install the prerequisite package ksh.x86_64, by issuing the yum command. For example:

```
yum install ksh.x86_64
```

Exception: You do not need to install the compat-libstdc++-33-3.2.3-69.el6.i686 and libstdc++.i686 libraries for Red Hat Enterprise Linux Version 7.1.

5. When the software installation is complete, you can restore the original YUM repository values by completing the following steps:

- a. Unmount the Red Hat Enterprise Linux installation DVD by issuing the following command:

```
umount /mnt
```

- b. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

- c. Rename the repo file that you created:

```
mv rhel71_dvd.repo rhel71_dvd.repo.orig
```

- d. Rename the original file to the original name:

```
mv rhel-source.repo.orig rhel-source.repo
```

6. Determine whether kernel parameter changes are required. Complete the following steps:

- a. Use the sysctl -a command to list the parameter values.
- b. Analyze the results by using the guidelines in Table 1 to determine whether any changes are required.
- c. If changes are required, set the parameters in the /etc/sysctl.conf file. The file changes are applied when the system is started.

Tip: Automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings. On Linux, the DB2® database software automatically adjusts interprocess communication (IPC) kernel parameter values to the preferred settings. For more information about kernel parameter settings, search for Linux kernel parameters in the IBM DB2 Version 11.1 product documentation.

Table 1. Linux kernel parameter optimum settings

Parameter	Description
-----------	-------------

Parameter	Description
kernel.shmmni	The maximum number of segments.
kernel.shmmax	The maximum size of a shared memory segment (bytes). This parameter must be set before automatically starting the IBM Spectrum Protect™ server on system startup.
kernel.shmall	The maximum allocation of shared memory pages (pages).
kernel.sem	(SEMMSL) The maximum semaphores per array.
There are four values for the kernel.sem parameter.	(SEMMNS) The maximum semaphores per system.
	(SEMOPM) The maximum operations per semaphore call.
	(SEMMNI) The maximum number of arrays.
kernel.msgmni	The maximum number of system-wide message queues.
kernel.msgmax	The maximum size of messages (bytes).
kernel.msgmnb	The default maximum size of queue (bytes).
kernel.randomize_va_space	The kernel.randomize_va_space parameter configures the use of memory ASLR for the kernel. Disable ASLR because it can cause errors for the DB2 software. To learn more details about the Linux ASLR and DB2, see technote 1365583.
vm.swappiness	The vm.swappiness parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the DB2 product information.
vm.overcommit_memory	The vm.overcommit_memory parameter influences how much virtual memory the kernel permits allocating. For more information about kernel parameters, see the DB2 product information.

7. Open firewall ports to communicate with the server. Complete the following steps:

- a. Determine the zone that is used by the network interface. The zone is public, by default.

Issue the following command:

```
# firewall-cmd --get-active-zones
public
  interfaces: ens4f0
```

- b. To use the default port address for communications with the server, open TCP/IP port 1500 in the Linux firewall.

Issue the following command:

```
firewall-cmd --zone=public --add-port=1500/tcp --permanent
```

If you want to use a value other than the default, you can specify a number in the range 1024 - 32767. If you open a port other than the default, you will need to specify that port when you run the configuration script.

- c. If you plan to use this system as a hub, open port 11090, which is the default port for secure (https) communications.

Issue the following command:

```
firewall-cmd --zone=public --add-port=11090/tcp --permanent
```

- d. Reload the firewall definitions for the changes to take effect.

Issue the following command:

```
firewall-cmd --reload
```

- Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 2. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 2. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf
Maximum number of open files	nofile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Windows systems

Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect™ server.

Procedure

- Install Windows Server 2016 Standard Edition, according to the manufacturer instructions.
- Change the Windows account control policies by completing the following steps.
 - Open the Local Security Policy editor by running `secpol.msc`.
 - Click Local Policies > Security Options and ensure that the following User Account Control policies are disabled:
 - Admin Approval Mode for the Built-in Administrator account
 - Run all administrators in Admin Approval Mode
- Configure your TCP/IP settings according to installation instructions for the operating system.
- Apply Windows updates and enable optional features by completing the following steps:
 - Apply the latest Windows Server 2016 updates.
 - Install and enable the Windows 2012 R2 feature Microsoft .NET Framework 3.5 from the Windows Server Manager.
 - If required, update the FC and Ethernet HBA device drivers to newer levels.
 - Install the multipath I/O driver that is appropriate for the disk system that you are using.
- Open the default TCP/IP port, 1500, for communications with the IBM Spectrum Protect server. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Backup server port 1500"  
dir=in action=allow protocol=TCP localport=1500
```

- On the Operations Center hub server, open the default port for secure (https) communications with the Operations Center. The port number is 11090. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Operations Center port 11090"  
dir=in action=allow protocol=TCP localport=11090
```

Configuring multipath I/O

You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.

- AIX systems
- Linux systems
- Windows systems

AIX systems

Procedure

1. Determine the Fibre Channel port address that you must use for the host definition on the disk subsystem. Issue the `lscfg` command for every port.

- o On small and medium systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
```

- o On large systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
lscfg -vps -l fcs2 | grep "Network Address"
lscfg -vps -l fcs3 | grep "Network Address"
```

2. Ensure that the following AIX® file sets are installed:

- o `devices.common.IBM.mpio.rte`
- o `devices.fcp.disk.array.rte`
- o `devices.fcp.disk.rte`

3. Issue the `cfgmgr` command to have AIX rescan the hardware and discover available disks. For example:

```
cfgmgr
```

4. To list the available disks, issue the following command:

```
lsdev -Ccdisk
```

You should see output similar to the following:

```
hdisk0 Available 00-00-00 SAS Disk Drive
hdisk1 Available 00-00-00 SAS Disk Drive
hdisk2 Available 01-00-00 SAS Disk Drive
hdisk3 Available 01-00-00 SAS Disk Drive
hdisk4 Available 06-01-02 MPIO IBM 2076 FC Disk
hdisk5 Available 07-01-02 MPIO IBM 2076 FC Disk
...
```

5. Use the output from the `lsdev` command to identify and list device IDs for each disk device.

For example, a device ID could be `hdisk4`. Save the list of device IDs to use when you create file systems for the IBM Spectrum Protect™ server.

6. Correlate the SCSI device IDs to specific disk LUNs from the disk system by listing detailed information about all physical volumes in the system. Issue the following command:

```
lspv -u
```

On an IBM® Storwize® system, the following information is an example of what is shown for each device:

```
hdisk4 00f8cf083fd97327 None active
3321360050763008101057800000000000003004214503IBMfcp
```

In the example, `60050763008101057800000000000030` is the UID for the volume, as reported by the Storwize management interface.

To verify disk size in megabytes and compare the value with what is listed for the system, issue the following command:

```
bootinfo -s hdisk4
```

Linux systems

Procedure

1. Edit the `/etc/multipath.conf` file to enable multipathing for Linux hosts. If the `multipath.conf` file does not exist, you can create it by issuing the following command:

```
mpathconf --enable
```

The following parameters were set in `multipath.conf` for testing on an IBM Storwize® system:

```
defaults {
    user_friendly_names no
}

devices {
    device {
        vendor "IBM "
        product "2145"
        path_grouping_policy group_by_prio
        user_friendly_names no
        path_selector "round-robin 0"
        prio "alua"
        path_checker "tur"
        failback "immediate"
        no_path_retry 5
        rr_weight uniform
        rr_min_io_rq "1"
        dev_loss_tmo 120
    }
}
```

2. Set the multipath option to start when the system is started. Issue the following commands:

```
systemctl enable multipathd.service
systemctl start multipathd.service
```

3. To verify that disks are visible to the operating system and are managed by multipath, issue the following command:

```
multipath -l
```

4. Ensure that each device is listed and that it has as many paths as you expect. You can use size and device ID information to identify which disks are listed.

For example, the following output shows that a 2 TB disk has two path groups and four active paths. The 2 TB size confirms that the disk corresponds to a pool file system. Use part of the long device ID number (12, in this example) to search for the volume on the disk-system management interface.

```
[root@tapsrv01 code]# multipath -l
36005076802810c50980000000000012 dm-43 IBM,2145
size=2.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=0 status=active
| |- 2:0:1:18 sdcw 70:64 active undef running
| `-- 4:0:0:18 sdgb 131:112 active undef running
`+- policy='round-robin 0' prio=0 status=enabled
|- 1:0:1:18 sdat 66:208 active undef running
`- 3:0:0:18 sddy 128:0 active undef running
```

- a. If needed, correct disk LUN host assignments and force a bus rescan. For example:

```
echo "- - -" > /sys/class/scsi_host/host0/scan
echo "- - -" > /sys/class/scsi_host/host1/scan
echo "- - -" > /sys/class/scsi_host/host2/scan
```

You can also restart the system to rescan disk LUN host assignments.

- b. Confirm that disks are now available for multipath I/O by reissuing the `multipath -l` command.

5. Use the multipath output to identify and list device IDs for each disk device.

For example, the device ID for your 2 TB disk is `36005076802810c50980000000000012`.

Save the list of device IDs to use in the next step.

Windows systems

Procedure

1. Ensure that the Multipath I/O feature is installed. If needed, install additional vendor-specific multipath drivers.
2. To verify that disks are visible to the operating system and are managed by multipath I/O, issue the following command:

```
c:\program files\IBM\SDDDSM\datapath.exe query device
```

3. Review the multipath output and ensure that each device is listed and that it has as many paths as you expect. You can use size and device serial information to identify which disks are listed.

For example, by using part of the long device serial number (34, in this example) you can search for the volume on the disk-system management interface. The 2 TB size confirms that the disk corresponds to a storage pool file system.

```
DEV#: 4 DEVICE NAME: Disk5 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 60050763008101057800000000000034 LUN SIZE: 2.0TB
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 0 0
1 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 27176 0
2 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 28494 0
3 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 0 0
```

4. Create a list of disk device IDs by using the serial numbers that are returned from the multipath output in the previous step.

For example, the device ID for your 2 TB disk is 60050763008101057800000000000034

Save the list of device IDs to use in the next step.

5. To bring new disks online and clear the read-only attribute, run diskpart.exe with the following commands. Repeat for each of the disks:

```
diskpart
select Disk 1
online disk
attribute disk clear readonly
select Disk 2
online disk
attribute disk clear readonly
< ... >
select Disk 49
online disk
attribute disk clear readonly
exit
```

Creating the user ID for the server

Create the user ID that owns the IBM Spectrum Protect™ server instance. You specify this user ID when you create the server instance during initial configuration of the server.

About this task

You can specify only lowercase letters (a-z), numerals (0-9), and the underscore character (_) for the user ID. The user ID and group name must comply with the following rules:

- The length must be 8 characters or fewer.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

Procedure

1. Use operating system commands to create a user ID.
 - o **AIX** | **Linux** Create a group and user ID in the home directory of the user that owns the server instance.

For example, to create the user ID `tsminst1` in group `tsmsrvrs` with a password of `tsminst1`, issue the following commands from an administrative user ID:

AIX

```
mkgroup id=1001 tsmsrvrs
mkuser id=1002 pgrp=tsmsrvrs home=/home/tsminst1 tsminst1
passwd tsminst1
```

Linux

```
groupadd tsmsrvrs
useradd -d /home/tsminst1 -m -g tsmsrvrs -s /bin/bash tsminst1
passwd tsminst1
```

Log off, and then log in to your system. Change to the user account that you created. Use an interactive login program, such as telnet, so that you are prompted for the password and can change it if necessary.

- o **Windows** Create a user ID and then add the new ID to the Administrators group. For example, to create the user ID tsminst1, issue the following command:

```
net user tsminst1 * /add
```

After you create and verify a password for the new user, add the user ID to the Administrators group by issuing the following commands:

```
net localgroup Administrators tsminst1 /add
net localgroup DB2ADMNS tsminst1 /add
```

2. Log off the new user ID.

Preparing file systems for the server

You must complete file system configuration for the disk storage to be used by the server.

- Preparing file systems on AIX systems
You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.
- Preparing file systems on Linux systems
You must format ext4 or xfs file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.
- Preparing file systems on Windows systems
You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.

Preparing file systems on AIX systems

You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.

Procedure

1. Increase the queue depth and maximum transfer size for all of the available *hdiskX* disks. Issue the following commands for each disk:

```
chdev -l hdisk4 -a max_transfer=0x100000
chdev -l hdisk4 -a queue_depth=32
chdev -l hdisk4 -a reserve_policy=no_reserve
chdev -l hdisk4 -a algorithm=round_robin
```

Do not run these commands for operating system internal disks, for example, *hdisk0*.

2. Create volume groups for the IBM Spectrum Protect™ database, active log, archive log, database backup, and storage pool. Issue the `mkvg` command, specifying the device IDs for corresponding disks that you previously identified.

For example, if the device names *hdisk4*, *hdisk5*, and *hdisk6* correspond to database disks, include them in the database volume group and so on.

System size: The following commands are based on the medium system configuration. For small and large systems, you must adjust the syntax as required.

```
mkvg -S -y tsmdb hdisk2 hdisk3 hdisk4
mkvg -S -y tsmactlog hdisk5
mkvg -S -y tsmarchlog hdisk6
mkvg -S -y tsmdbback hdisk7 hdisk8 hdisk9 hdisk10
mkvg -S -y tsmstgpool hdisk11 hdisk12 hdisk13 hdisk14 ... hdisk49
```

3. Determine the physical volume names and the number of free physical partitions to use when you create logical volumes. Issue the `lsvg` for each volume group that you created in the previous step.

For example:

```
lsvg -p tsmdb
```

The output is similar to the following. The *FREE PPs* column represents the free physical partitions:

```
tsmdb:
PV_NAME  PV STATE  TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk4   active    1631       1631      327..326..326..326..326
hdisk5   active    1631       1631      327..326..326..326..326
hdisk6   active    1631       1631      327..326..326..326..326
```

4. Create logical volumes in each volume group by using the `mklv` command. The volume size, volume group, and device names vary, depending on the size of your system and variations in your disk configuration.

For example, to create the volumes for the IBM Spectrum Protect database on a medium system, issue the following commands:

```
mklv -y tsmdb00 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk2
mklv -y tsmdb01 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk3
mklv -y tsmdb02 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk4
```

5. Format file systems in each logical volume by using the `crfs` command.

For example, to format file systems for the database on a medium system, issue the following commands:

```
crfs -v jfs2 -d tsmdb00 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace00 -A yes
crfs -v jfs2 -d tsmdb01 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace01 -A yes
crfs -v jfs2 -d tsmdb02 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace02 -A yes
```

6. Mount all of the newly created file systems by issuing the following command:

```
mount -a
```

7. List all file systems by issuing the `df` command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example of command output shows that the amount of used space is typically 1%:

```
tapsrv07> df -g /tsminst1/*
Filesystem      GB blocks  Free    %Used  Iused  %Iused  Mounted on
/dev/tsmact00   195.12    194.59  1%      4      1%      /tsminst1/TSMalog
```

8. Verify that the user ID you created in *Creating the user ID for the server* has read and write access to the directories for the server.

Preparing file systems on Linux systems

You must format `ext4` or `xfs` file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Using the list of device IDs that you generated previously, issue the `mkfs` command to create and format a file system for each storage LUN device. Specify the device ID in the command. See the following examples. For the database, format `ext4` file systems:

```
mkfs -t ext4 -T largefile -m 2 /dev/mapper/36005076802810c50980000000000012
```

For storage pool LUNs, format `xfs` file systems:

```
mkfs -t xfs /dev/mapper/3600507630081010578000000000002c3
```

You might issue the `mkfs` command as many as 50 times, depending on how many different devices you have.

2. Create mount point directories for file systems.

Issue the `mkdir` command for each directory that you must create. Use the directory values that you recorded in the planning worksheets.

For example, to create the server instance directory by using the default value, issue the following command:

```
mkdir /tsminst1
```


Repeat the mkdir command for each file system.

3. Add an entry in the /etc/fstab file for each file system so that file systems are mounted automatically when the server is started.

For example:

```
/dev/mapper/36005076802810c50980000000000012 /tsminst1/TSMdbspace00 ext4 defaults 0 0
```

4. Mount the file systems that you added to the /etc/fstab file by issuing the mount -a command.
5. List all file systems by issuing the df command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example on an IBM® Storwize® system shows that the amount of used space is typically 1%:

```
[root@tapsrv04 ~]# df -h /tsminst1/*
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/36005076300810105780000000000003 134G  188M 132G   1% /tsminst1/TSMalog
```

6. Verify that the user ID you created in Creating the user ID for the server has read and write access to the directories for the IBM Spectrum Protect server.

Preparing file systems on Windows systems

You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Create mount point directories for file systems.

Issue the md command for each directory that you must create. Use the directory values that you recorded in the planning worksheets. For example, to create the server instance directory by using the default value, issue the following command:

```
md c:\tsminst1
```

Repeat the md command for each file system.

2. Create a volume for every disk LUN that is mapped to a directory under the server instance directory by using the Windows volume manager.

Go to Server Manager > File and Storage Services and complete the following steps for each disk that corresponds to the LUN mapping that was created in the previous step:

- a. Bring the disk online.
- b. Initialize the disk to the GPT basic type, which is the default.
- c. Create a simple volume that occupies all of the space on the disk. Format the file system by using NTFS, and assign a label that matches the purpose of the volume, such as TSMfile00. Do not assign the new volume to a drive letter. Instead, map the volume to a directory under the instance directory, such as C:\tsminst1\TSMfile00.

Tip: Determine the volume label and directory mapping labels based on the size of the disk that is reported.

3. Verify that file systems are mounted at the correct LUN and correct mount point. List all file systems by issuing the mountvol command and then review the output. For example:

```
\\?\Volume{8ffb9678-3216-474c-a021-20e420816a92}\
C:\tsminst1\TSMdbspace00\
```

4. After the disk configuration is complete, restart the system.

What to do next

You can confirm the amount of free space for each volume by using Windows Explorer.

Installing the server and Operations Center

Use the IBM® Installation Manager graphical wizard to install the components.

- Installing on AIX and Linux systems
Install the IBM Spectrum Protect™ server and the Operations Center on the same system.

- Installing on Windows systems
Install the IBM Spectrum Protect server and the Operations Center on the same system.

Installing on AIX® and Linux systems

Install the IBM Spectrum Protect™ server and the Operations Center on the same system.

Before you begin

Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

1. **AIX** Verify that the required RPM files are installed on your system.

See Installing prerequisite RPM files for the graphical wizard for details.

2. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042992.
3. Go to Passport Advantage® and download the package file to an empty directory of your choice.
4. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

5. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file.

6. **AIX** Ensure that the following command is enabled so that the wizards work properly:

```
lsuser
```

By default, the command is enabled.

7. Change to the directory where you placed the executable file.
8. Start the installation wizard by issuing the following command:

```
./install.sh
```

When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.
- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.
- Installing prerequisite RPM files for the graphical wizard
RPM files are required for the IBM Installation Manager graphical wizard.

Related tasks:

- [Other methods for installing IBM Spectrum Protect components \(AIX\)](#)
- [Other methods for installing IBM Spectrum Protect components \(Linux\)](#)

Installing on Windows systems

Install the IBM Spectrum Protect™ server and the Operations Center on the same system.

Before you begin

Make sure that the following prerequisites are met:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

1. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042993.
2. Go to Passport Advantage® and download the package file to an empty directory of your choice.
3. Change to the directory where you placed the executable file.
4. Double-click the executable file to extract to the current directory.
5. In the directory where the installation files were extracted, start the installation wizard by double-clicking the install.bat file. When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.

- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.

Related tasks:

 Other methods for installing IBM Spectrum Protect components

Configuring the server and the Operations Center

After you install the components, complete the configuration for the IBM Spectrum Protect™ server and the Operations Center.

- **Configuring the server instance**
Use the IBM Spectrum Protect server instance configuration wizard to complete the initial configuration of the server.
- **Installing the backup-archive client**
As a best practice, install the IBM Spectrum Protect backup-archive client on the server system so that the administrative command-line client and scheduler are available.
- **Setting options for the server**
Review the server options file that is installed with the IBM Spectrum Protect server to verify that the correct values are set for your system.
- **Configuring secure communications with Transport Layer Security**
To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) is enabled on the IBM Spectrum Protect server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.
- **Configuring the Operations Center**
After you install the Operations Center, complete the following configuration steps to start managing your storage environment.
- **Registering the product license**
To register your license for the IBM Spectrum Protect product, use the REGISTER LICENSE command.
- **Configuring data deduplication**
Create a directory-container storage pool and at least one directory to use inline data deduplication.
- **Defining data retention rules for your business**
After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.
- **Defining schedules for server maintenance activities**
Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations Center command builder.

- Defining client schedules
Use the Operations Center to create schedules for client operations.

Configuring the server instance

Use the IBM Spectrum Protect™ server instance configuration wizard to complete the initial configuration of the server.

Before you begin

Ensure that the following requirements are met:

AIX | **Linux**

- The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/` directory. Also, ensure that the SSH daemon service has access rights to connect to the system by using the `localhost` value.
- You must be able to log in to IBM Spectrum Protect with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

Windows

Verify that the remote registry service is started by completing the following steps:

1. Click Start > Administrative Tools > Services. In the Services window, select Remote Registry. If it is not started, click Start.
2. Ensure that port 137, 139, and 445 are not blocked by a firewall:
 - a. Click Start > Control Panel > Windows Firewall.
 - b. Select Advanced Settings.
 - c. Select Inbound Rules.
 - d. Select New Rule.
 - e. Create a port rule for TCP ports 137, 139, and 445 to allow connections for domain and private networks.
3. Configure the user account control by accessing the local security policy options and completing the following steps.
 - a. Click Start > Administrative Tools > Local Security Policy. Expand Local Policies > Security Options.
 - b. If not already enabled, enable the built-in administrator account by selecting Accounts: Administrator account status > Enable > OK.
 - c. If not already disabled, disable user account control for all Windows administrators by selecting User Account Control: Run all administrators in Admin Approval Mode > Disable > OK.
 - d. If not already disabled, disable the User Account Control for the built-in Administrator account by selecting User Account Control: Admin Approval Mode for the Built-in Administrator Account > Disable > OK.
4. If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

About this task

The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Procedure

1. Start the local version of the wizard.
 - o **AIX** | **Linux** Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be only run as a root user.
 - o **Windows** Click Start > All Programs > IBM Spectrum Protect > Configuration Wizard.
2. Follow the instructions to complete the configuration. Use the information that you recorded in Planning worksheets during IBM Spectrum Protect system set up to specify directories and options in the wizard.

AIX | **Linux**

On the Server Information window, set the server to start automatically by using the instance user ID when the system boots.

Windows

By using the configuration wizard, the server is set to start automatically when rebooted.

Installing the backup-archive client

As a best practice, install the IBM Spectrum Protect™ backup-archive client on the server system so that the administrative command-line client and scheduler are available.

Procedure

To install the backup-archive client, follow the installation instructions for your operating system.

- Install UNIX and Linux backup-archive clients
- Installing the Windows client for the first time

Setting options for the server

Review the server options file that is installed with the IBM Spectrum Protect™ server to verify that the correct values are set for your system.

Procedure

1. Go to the server instance directory and open the dsmserv.opt file.
2. Review the values in the following table and verify your server option settings, based on system size.

Server option	Small system value	Medium system value	Large system value
ACTIVELOGDIRECTORY	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
ACTIVELOGSIZE	131072	131072	262144
ARCHLOGCOMPRESS	Yes	No	No
ARCHLOGDIRECTORY	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
COMMMETHOD	TCPIP	TCPIP	TCPIP
COMMTIMEOUT	3600	3600	3600
DEDUPREQUIRESBACKUP	No	No	No
DEVCONFIG	devconf.dat	devconf.dat	devconf.dat
EXPINTERVAL	0	0	0
IDLETIMEOUT	60	60	60
MAXSESSIONS	250	500	1000
NUMOPENVOLSALLOWED	20	20	20
TCPADMINPORT	1500	1500	1500
TCPPORT	1500	1500	1500
VOLUMEHISTORY	volhist.dat	volhist.dat	volhist.dat

Update server option settings if necessary, to match the values in the table. To make updates, close the dsmserv.opt file and use the SETOPT command from the administrative command-line interface to set the options.

For example, to update the IDLETIMEOUT option to 60, issue the following command:

```
setopt idletimeout 60
```

3. To configure secure communications for the server, clients, and the Operations Center, verify the options in the following table.

Server option	All system sizes
SSLFIPSMODE	NO
TCPPORT	Specify the port number on which the server waits for requests for TCP/IP and SSL-enabled sessions from the client.

Server option	All system sizes
TCPADMINPORT	Specify the port address on which the server waits for requests for TCP/IP and SSL-enabled sessions from the command-line administrative client.

If any of the option values must be updated, edit the dsmserv.opt file by using the following guidelines:

- o Remove the asterisk at the beginning of a line to enable an option.
- o On each line, enter only one option and the specified value for the option.
- o If an option occurs in multiple entries in the file, the server uses the last entry.

Save your changes and close the file. If you edit the dsmserv.opt file directly, you must restart the server for the changes to take effect.

Related reference:

- 🔗 [Server options reference](#)
- 🔗 [SETOPT \(Set a server option for dynamic update\)](#)

Configuring secure communications with Transport Layer Security

To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) is enabled on the IBM Spectrum Protect™ server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.

About this task

As shown in the following figure, you can manually configure secure communications between the server and backup-archive client by setting options in the server and client options files, and then transferring the self-signed certificate that is generated on the server to the client. Alternatively, you can obtain and transfer a unique certificate that is signed by a certificate authority (CA).



For more information about configuring the server and clients for SSL or TLS communications, see [Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL](#).

Configuring the Operations Center

After you install the Operations Center, complete the following configuration steps to start managing your storage environment.

Before you begin

When you connect to the Operations Center for the first time, you must provide the following information:

- Connection information for the server that you want to designate as a hub server
- Login credentials for an administrator ID that is defined for that server

Procedure

1. Designate the hub server. In a web browser, enter the following address:

```
https://hostname:secure_port/oc
```

where:

- o *hostname* represents the name of the computer where the Operations Center is installed
- o *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer

For example, if your host name is `tsm.storage.mylocation.com` and you are using the default secure port for the Operations Center, which is 11090, the address is:

```
https://tsm.storage.mylocation.com:11090/oc
```

When you log in to the Operations Center for the first time, a wizard guides you through an initial configuration to set up a new administrator with system authority on the server.

2. Set up secure communications between the Operations Center and the hub server by configuring the Secure Sockets Layer (SSL) protocol.

Follow the instructions in [Securing communications between the Operations Center and the hub server](#).

3. Optional: To receive a daily email report that summarizes system status, configure your email settings in the Operations Center.

Follow the instructions in [Tracking system status by using email reports](#).

- [Securing communications between the Operations Center and the hub server](#)
To secure communications between the Operations Center and the hub server, add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.

Registering the product license


To register your license for the IBM Spectrum Protect™ product, use the REGISTER LICENSE command.

About this task

Licenses are stored in enrollment certificate files, which contain licensing information for the product. The enrollment certificate files are on the installation media, and are placed on the server during installation. When you register the product, the licenses are stored in a NODELOCK file within the current directory.

Procedure


Register a license by specifying the name of the enrollment certificate file that contains the license. To use the Operations Center command builder for this task, complete the following steps.

1. Open the Operations Center.
2. Open the Operations Center command builder by hovering over the settings icon  and clicking Command Builder.
3. Issue the REGISTER LICENSE command. For example, to register a base IBM Spectrum Protect license, issue the following command:

```
register license file=tsmbasic.lic
```

What to do next

Save the installation media that contains your enrollment certificate files. You might need to register your license again if, for example, one of the following conditions occur:

- The server is moved to a different computer.
- The NODELOCK file is corrupted. The server stores license information in the NODELOCK file, which is in the directory from which the server is started.
-  If you change the processor chip that is associated with the server on which the server is installed.

Related reference:

[REGISTER LICENSE \(Register a new license\)](#)

Configuring data deduplication

Create a directory-container storage pool and at least one directory to use inline data deduplication.

Before you begin

Use the storage pool directory information that you recorded in Planning worksheets for this task.

Procedure

1. Open the Operations Center.
2. On the Operations Center menu bar, hover over Storage.
3. From the list that is displayed, click Storage Pools.
4. Click the +Storage Pool button.
5. Complete the steps in the Add Storage Pool wizard:
 - o To use inline data deduplication, select a Directory storage pool under Container-based storage.
 - o When you configure directories for the directory-container storage pool, specify the directory paths that you created for storage during system setup.
6. After you configure the new directory-container storage pool, click Close & View Policies to update a management class and start using the storage pool.

Defining data retention rules for your business

After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.

Procedure

1. On the Services page of the Operations Center, select the STANDARD domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab. The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.
3. Click the Configure toggle, and make the following changes:
 - o Change the backup destination for the STANDARD management class to the directory-container storage pool.
 - o Change the value for the Backups column to No limit.
 - o Change the retention period. Set the Keep Extra Backups column to 30 days or more, depending on your business requirements.
4. Save your changes and click the Configure toggle again so that the policy set is no longer editable.
5. Activate the policy set by clicking Activate.

Related tasks:

Specifying rules for backing up and archiving client data

Defining schedules for server maintenance activities

Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations Center command builder.

About this task

Schedule server maintenance operations to run after client backup operations. You can control the timing of schedules by setting the start time in combination with the duration time for each operation.

The following example shows how you can schedule server maintenance operations in combination with the client backup schedule for a single-site disk solution.

Operation	Schedule
Client backup	Starts at 22:00.
Processing for database and disaster recovery files	<ul style="list-style-type: none">• The database backup operation starts at 11:00, or 13 hours after the beginning of the client backup operation. This process runs until completion.• Device configuration information and volume history backup operations start at 17:00, or 6 hours after the start of the database backup operation.• Volume history deletion starts at 20:00, or 9 hours after the start of the database backup operation.
Inventory expiration	Starts at 12:00, or 14 hours after the beginning of the client backup operation. This process runs until completion.

Procedure

After you configure the device class for the database backup operations, create schedules for database backup and other required maintenance operations by using the DEFINE SCHEDULE command. Depending on the size of your environment, you might need to adjust the start times for each schedule in the example.

1. Define a device class for the backup operations. For example, use the DEFINE DEVCLASS command to create a device class that is named DBBACK_FILEDEV:

```
define devclass dbback_filedev devtype=file
  directory=db_backup_directories
```

where *db_backup_directories* is a list of the directories that you created for the database backup.

AIX | **Linux** For example, if you have four directories for database backups, starting with /tsminst1/TSMbkup00, issue the following command:

```
define devclass dbback_filedev devtype=file
  directory=/tsminst1/TSMbkup00,
  /tsminst1/TSMbkup01,/tsminst1/TSMbkup02,
  /tsminst1/TSMbkup03"
```

Windows For example, if you have four directories for database backups, starting with C:\tsminst1\TSMbkup00, issue the following command:

```
define devclass dbback_filedev devtype=file
  directory="c:\tsminst1\TSMbkup00,
  c:\tsminst1\TSMbkup01,c:\tsminst1\TSMbkup02,
  c:\tsminst1\TSMbkup03"
```

2. Set the device class for automatic database backup operations. Use the SET DBRECOVERY command to specify the device class that you created in the preceding step. For example, if the device class is dbback_filedev, issue the following command:

```
set dbrecovery dbback_filedev
```

3. Create schedules for the maintenance operations by using the DEFINE SCHEDULE command. See the following table for the required operations with examples of the commands.

Operation	Example command
Back up the database.	<p>Create a schedule to run the BACKUP DB command. If you are configuring a small system, set the COMPRESS parameter to YES. For example, on a small system, issue the following command to create a backup schedule that uses the new device class:</p> <pre>define schedule DBBACKUP type=admin cmd="backup db devclass=dbback_filedev type=full numstreams=3 wait=yes compress=yes" active=yes desc="Back up the database." startdate=today starttime=11:00:00 duration=45 durunits=minutes</pre>
Back up the device configuration information.	<p>Create a schedule to run the BACKUP DEVCONFIG command:</p> <pre>define schedule DEVCONFIGBKUP type=admin cmd="backup devconfig filenames=devconfig.dat" active=yes desc="Backup the device configuration file." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>
Back up the volume history.	<p>Create a schedule to run the BACKUP VOLHISTORY command:</p> <pre>define schedule VOLHISTBKUP type=admin cmd="backup volhistory filenames=volhist.dat" active=yes desc="Back up the volume history." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>

Operation	Example command
Remove older versions of database backups that are no longer required.	<p>Create a schedule to run the DELETE VOLHISTORY command:</p> <pre>define schedule DELVOLHIST type=admin cmd="delete volhistory type=dbb todate=today-6 totime=now" active=yes desc="Remove old database backups." startdate=today starttime=20:00:00 duration=45 durunits=minutes</pre>
Remove objects that exceed their allowed retention.	<p>Create a schedule to run the EXPIRE INVENTORY command. Set the RESOURCE parameter based on the system size that you are configuring:</p> <ul style="list-style-type: none"> o Small systems: 10 o Medium systems: 30 o Large systems: 40 <p>For example, on a medium-sized system, issue the following command to create a schedule that is named EXPINVENTORY:</p> <pre>define schedule EXPINVENTORY type=admin cmd="expire inventory wait=yes resource=30 duration=120" active=yes desc="Remove expired objects." startdate=today starttime=12:00:00 duration=45 durunits=minutes</pre>

What to do next

After you create schedules for the server maintenance tasks, you can view them in the Operations Center by completing the following steps:

1. On the Operations Center menu bar, hover over Servers.
2. Click Maintenance.

Related reference:

[DEFINE SCHEDULE](#) (Define a schedule for an administrative command)

Defining client schedules

Use the Operations Center to create schedules for client operations.

Procedure

1. On the Operations Center menu bar, hover over Clients.
2. Click Schedules.
3. Click +Schedule.
4. Complete the steps in the Create Schedule wizard. Set client backup schedules to start at 22:00, based on the server maintenance activities that you scheduled in Defining schedules for server maintenance activities.

Installing and configuring backup-archive clients

Following the successful setup of your IBM Spectrum Protect™ server system, install and configure client software to begin backing up data.

Procedure

To install the backup-archive client, follow the installation instructions for your operating system.

- Install UNIX and Linux backup-archive clients
- Installing the Windows client for the first time

What to do next

Register and assign your clients to schedules.

- Registering and assigning clients to schedules
Add and register your clients through the Operations Center by using the Add Client wizard.
- Installing the client management service
Install the client management service for backup-archive clients that run on Linux and Windows operating systems. The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Registering and assigning clients to schedules

Add and register your clients through the Operations Center by using the Add Client wizard.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see technote 7048963.

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see Authenticating users by using an Active Directory database.

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the REGISTER NODE command and specify the USERID parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see Register a node.

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard.

Complete the following steps:

- a. On the Operations Center menu bar, click Clients.
- b. In the Clients table, click +Client.
- c. Complete the steps in the Add Client wizard:
 - i. Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the Enable check box.
 - ii. In the Configuration window, copy the TCPSERVERADDRESS, TCPPORT, NODENAME, and DEDUPLICATION option values.
Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
 - iii. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
 - iv. Set how risks are displayed for the client by specifying the at-risk setting.
 - v. Click Add Client.

Installing the client management service

Install the client management service for backup-archive clients that run on Linux and Windows operating systems. The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Procedure

Install the client management service on the same computer as the backup-archive client by completing the following steps:

1. Download the installation package for the client management service from an IBM® download site such as IBM Passport Advantage® or IBM Fix Central. Look for a file name that is similar to `<version>-IBM_Spectrum_Protect-CMS-`

`operating_system.bin`.

2. Create a directory on the client system that you want to manage, and copy the installation package there.
 3. Extract the contents of the installation package file.
 4. Run the installation batch file from the directory where you extracted the installation and associated files. This is the directory that you created in step 2.
 5. To install the client management service, follow the instructions in the IBM Installation Manager wizard. If IBM Installation Manager is not already installed on the client system, you must select both IBM Installation Manager and IBM Spectrum Protect™ Client Management Services.
- Verifying that the client management service is installed correctly
Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.
 - Configuring the Operations Center to use the client management service
If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Related tasks:

[↗](#) Configuring the client management service for custom client installations

Verifying that the client management service is installed correctly

Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.

Procedure

On the client system, at the command line, run the following commands to view the configuration of the client management service:

- On Linux client systems, issue the following command:

```
client_install_dir/cms/bin/CmsConfig.sh list
```

where `client_install_dir` is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
/opt/tivoli/tsm/cms/bin/CmsConfig.sh list
```

The output is similar to the following text:

Listing CMS configuration

```
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys

  Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

- On Windows client systems, issue the following command:

```
client_install_dir\cms\bin\CmsConfig.bat list
```

where `client_install_dir` is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
C:\Program Files\Tivoli\TSM\cms\bin\CmsConfig.bat list
```

The output is similar to the following text:

Listing CMS configuration

```
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt
```

```
Log File: C:\Program Files\Tivoli\TSM\baclient\dserror.log
en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

```
Log File: C:\Program Files\Tivoli\TSM\baclient\dmsched.log
en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

If the client management service is correctly installed and configured, the output displays the location of the error log file. The output text is extracted from the following configuration file:

- On Linux client systems:

```
client_install_dir/cms/Liberty/usr/servers/cmsServer/client-configuration.xml
```

- On Windows client systems:

```
client_install_dir\cms\Liberty\usr\servers\cmsServer\client-configuration.xml
```

If the output does not contain any entries, you must configure the client-configuration.xml file. For instructions to configure this file, see [Configuring the client management service for custom client installations](#). You can use the `CmsConfig verify` command to verify that a node definition is correctly created in the client-configuration.xml file.

Configuring the Operations Center to use the client management service

If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Before you begin

Ensure that the client management service is installed and started on the client system. Verify whether the default configuration is used. The default configuration is not used if either of the following conditions is met:

- The client management service does not use the default port number, 9028.
- The backup-archive client is not accessed by the same IP address as the client system where the backup-archive client is installed. For example, a different IP address might be used in the following situations:
 - The computer system has two network cards. The backup-archive client is configured to communicate on one network, while the client management service communicates on the other network.
 - The client system is configured with the Dynamic Host Configuration Protocol (DHCP). As a result, the client system is dynamically assigned an IP address, which is saved on the server during the previous backup-archive client operation. When the client system is restarted, the client system might be assigned a different IP address. To ensure that the Operations Center can always find the client system, you specify a fully qualified domain name.

Procedure

To configure the Operations Center to use the client management service, complete the following steps:

1. On the **Clients** page of the Operations Center, select the client.
2. Click **Details > Properties**.
3. In the **Remote diagnostics URL** field in the **General** section, specify the URL for the client management service on the client system. The address must start with `https`. The following table shows examples of the remote diagnostics URL.

Type of URL	Example
With DNS host name and default port, 9028	<code>https://server.example.com</code>
With DNS host name and non-default port	<code>https://server.example.com:1599</code>
With IP address and non-default port	<code>https://192.0.2.0:1599</code>

4. Click **Save**.

What to do next

You can access client diagnostic information such as client log files from the **Diagnosis** tab in the Operations Center.

Completing the implementation

After the IBM Spectrum Protect™ solution is configured and running, test backup operations and set up monitoring to ensure that everything runs smoothly.

Procedure

1. Test backup operations to verify that your data is protected in the way that you expect.
 - a. On the Clients page of the Operations Center, select the clients that you want to back up, and click Back Up.
 - b. On the Servers page of the Operations Center, select the server for which you want to back up the database. Click Back Up, and follow the instructions in the Back Up Database window.
 - c. Verify that the backup operations completed successfully with no warning or error messages.
Tip: Alternatively, you can use the backup-archive client GUI to back up client data and you can backup the server database by issuing BACKUP DB command from an administrative command-line.
2. Set up monitoring for your solution by following the instructions in Monitoring a single-site disk solution.

Monitoring a single-site disk solution

After you implement a single-site disk solution with IBM Spectrum Protect™, monitor the solution for correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

About this task

The preferred way to monitor a solution is by using the Operations Center, which provides overall and detailed system status in a graphical user interface. In addition, you can configure the Operations Center to generate a daily email report that summarizes system status.

In some cases, you might want to use advanced monitoring tools to complete specific monitoring or troubleshooting tasks.

Tip: If you plan to diagnose issues with backup-archive clients on Linux or Windows operating systems, install IBM Spectrum Protect client management services on each computer where a backup-archive client is installed. In this way, you can ensure that the Diagnose button is available in the Operations Center for diagnosing issues with backup-archive clients. To install the client management service, follow the instructions in Installing the client management service.

Procedure

1. Complete daily monitoring tasks. For instructions, see Daily monitoring checklist.
2. Complete periodic monitoring tasks. For instructions, see Periodic monitoring checklist.
3. To verify that your IBM Spectrum Protect solution complies with licensing requirements, follow the instructions in Verify license compliance.
4. To set up Operations Center to generate email status reports, see Tracking system status by using email reports

What to do next

Resolve any issues that you detect. To resolve an issue by changing the configuration of your solution, follow the instructions in Managing operations for a single-site disk solution. The following resources are also available:


- To resolve performance issues, see Performance.
- To resolve other types of issues, see Troubleshooting.


Daily monitoring checklist

To ensure that you are completing the daily monitoring tasks for your IBM Spectrum Protect™ solution, review the daily monitoring checklist.

Complete the daily monitoring tasks from the Operations Center Overview page. You can access the Overview page by opening the Operations Center and clicking Overviews.

The following figure shows the location for completing each task.








Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.


The following table lists the daily monitoring tasks and provides instructions for completing each task.

Table 1. Daily monitoring tasks



Task	Basic procedures	Advanced procedures and troubleshooting information
------	------------------	---



Task	Basic procedures	Advanced procedures and troubleshooting information
<p>Watch for security notifications, which can indicate a ransomware attack.</p>	<p>If a potential ransomware attack is detected in the IBM Spectrum Protect environment, a security notification message is displayed in the foreground of the Operations Center. For more information, click the message to open the Security Notifications page.</p>	<p>On the Security Notifications page, you can take the following actions:</p> <ul style="list-style-type: none"> • View notification details by client. Restriction: In Operations Center Version 8.1.5, notifications are available only for backup-archive clients. • Acknowledge a security notification by selecting it and clicking Acknowledge. When you acknowledge a security notification, a check mark is added to the Acknowledged column of the Security Notifications page for the selected client. The standard by which a notification is acknowledged is determined by your organization. A check mark might mean that you investigated the issue and determined that it is a false positive. Or it might mean that a problem exists and is being resolved. • Assign a security notification to an administrator by selecting the security notification and clicking Assign. To view the assignment, the administrator must sign in to the Operations Center and click Overviews > Security. If you are not certain whether the administrator regularly monitors the Security Notifications page, notify the administrator about the assignment. • If the notification is a false positive, you can select the security notification and click Reset. The security notification is deleted. Historical data that is used for baseline comparisons with the most recent backup operation is deleted. A new baseline is calculated going forward.
<p>1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.</p>	<p>To verify whether clients are at risk, in the Clients area, look for an At risk notification. To view details, click the Clients area. Attention: If the At risk percentage is much greater than usual, it might indicate a ransomware attack. A ransomware attack can cause backup operations to fail, thus placing clients at risk. For example, if the percentage of clients at risk is normally between 5% and 10%, but the percentage increases to 40% or 50%, investigate the cause. If you installed the client management service on a backup-archive client, you can view and analyze the client error and schedule logs by completing the following steps:</p> <ol style="list-style-type: none"> 1. In the Clients table, select the client and click Details. 2. To diagnose an issue, click Diagnosis. 	<p>For clients that do not have the client management service installed, access the client system to review the client error logs.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>2 Determine whether client-related or server-related errors require attention.</p>	<p>To determine the severity of any reported alerts, in the Alerts area, hover over the columns.</p>	<p>To view additional information about alerts, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Alerts area. 2. In the Alerts table, select an alert. 3. In the Activity Log pane, review the messages. The pane displays related messages that were issued before and after the selected alert occurred.
<p>3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.</p>	<ol style="list-style-type: none"> 1. To verify whether servers are at risk, in the Servers area, look for an Unavailable notification. 2. To view additional information, click the Servers area. 3. Select a server in the Servers table and click Details. 	<p>Tip: If you detect an issue that is related to server properties, update the server properties:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a server and click Details. 2. To update server properties, click Properties.
<p>4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.</p>	<ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Status column of the table, view the status of the server and resolve any issues: <ul style="list-style-type: none"> o Normal  Sufficient space is available for the server database, active log, and archive log. o Critical  Insufficient space is available for the server database, active log, or archive log. You must add space immediately, or the data protection services that are provided by the server will be interrupted. o Warning  The server database, active log, or archive log is running out of space. If this condition persists, you must add space. o Unavailable  Status cannot be obtained. Ensure that the server is running, and that there are no network issues. This status is also shown if the monitoring administrator ID is locked or otherwise unavailable on the server. This ID is named IBM-OC-hub_server_name. o Unmonitored  Unmonitored servers are defined to the hub server, but are not configured for management by the Operations Center. To configure an unmonitored server, select the server, and click Monitor Spoke. 	<p>You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>5 Verify server database backup operations.</p>	<p>To determine when a server was most recently backed up, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column. 	<p>To obtain more detailed information about backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a row and click Details. 2. In the DB Backup area, hover over the check marks to review information about backup operations. <p>If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click Back Up. <p>To determine whether the server database is configured for automatic backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the menu bar, hover over the settings icon  and click Command Builder. 2. Issue the QUERY DB command: <pre>query db f=d</pre> 3. In the output, review the Full Device Class Name field. If a device class is specified, the server is configured for automatic database backups.
<p>6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.</p>	<p>To search for information about processes that failed because of server issues, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click Servers > Maintenance. 2. To obtain the two-week history of a process, view the History column. 3. To obtain more information about a scheduled process, hover over the check box that is associated with the process. 	<p>For more information about monitoring processes and resolving issues, see the Operations Center online help.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>7 Verify that the amount of data that was recently sent to and from servers is within the expected range.</p>	<ul style="list-style-type: none"> • To obtain an overview of activity in the last 24 hours, view the Activity area. • To compare activity in the last 24 hours with activity in the previous 24 hours, review the figures in the Current and Previous areas. 	<ul style="list-style-type: none"> • If more data was sent to the server than you expected, determine which clients are backing up more data and investigate the cause. It is possible that client-side data deduplication is not working correctly. Attention: If the amount of backed-up data is significantly larger than usual, it might indicate a ransomware attack. When ransomware encrypts data, the system perceives the data as being changed, and the changed data is backed up. Thus, backup volumes become larger. To determine which clients are affected, click the Applications, Virtual, or Systems tabs. • If less data was sent to the server than you expected, investigate whether client backup operations are proceeding on schedule.

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>8 Verify that storage pools are available to back up client data.</p>	<ol style="list-style-type: none"> 1. If problems are indicated in the Storage & Data Availability area, click Pools to view the details: <ul style="list-style-type: none"> o If the Critical  status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. Attention: If the status is critical, investigate the cause: <ul style="list-style-type: none"> ■ If the data deduplication rate for a storage pool drops significantly, it might indicate a ransomware attack. During a ransomware attack, data is encrypted and cannot be deduplicated. To verify the data deduplication rate, in the Storage Pools table, review the value in the % Savings column. ■ If a storage pool unexpectedly becomes 100% utilized, it might indicate a ransomware attack. To verify the utilization, review the value in the Capacity Used column. Hover over the values to see the percentages of used and free space. o If the Warning  status is displayed, the storage pool is running out of space, or its access status is read-only. 2. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column. 	<p>To view the storage-pool capacity that was used over the past two weeks, select a row in the Storage Pools table and click Details.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>9 Verify that storage devices are available for backup operations.</p>	<p>In the Storage & Data Availability area, in the Volumes section, under the capacity bars, review the status that is reported next to Devices. If a Critical  or Warning  status is displayed for any device, investigate the issue. To view details, click Devices.</p>	<p>Disk devices might have a critical or warning status for the following reasons:</p> <ul style="list-style-type: none"> • For DISK device classes, volumes might be offline or have a read-only access status. The Disk Storage column of the Disk Devices table shows the state of volumes. • For FILE device classes that are not shared, directories might be offline. Also, insufficient free space might be available for allocating scratch volumes. The Disk Storage column of the Disk Devices table shows the state of directories. • For FILE device classes that are shared, drives might be unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. Other columns of the Disk Devices table show the state of the drives and paths.

Periodic monitoring checklist

To help ensure that your IBM Spectrum Protect™ solution operates correctly, complete the tasks in the periodic monitoring checklist. Schedule periodic tasks frequently enough so that you can detect potential issues before they become problematic.




Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.

Table 1. Periodic monitoring tasks


Task	Basic procedures	Advanced procedures and troubleshooting
------	------------------	---

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor system performance.	<p>Determine the length of time that is required for client backup operations:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. Find the server that is associated with the client. 2. Click Servers. Select the server and click Details. 3. To view the duration of completed tasks in the last 24 hours, click Completed Tasks. 4. To view the duration of tasks that were completed more than 24 hours ago, use the QUERY ACTLOG command. Follow the instructions in . 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. <p>If you installed the client management service on a backup-archive client, you can diagnose performance issues for the backup-archive client by completing the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. 2. Select a backup-archive client and click Details. 3. To retrieve client logs, click Diagnosis. 	<p>For instructions about reducing the time that it takes for the client to back up data to the server, see Resolving common client performance problems.</p> <p>Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks.</p> <p>For information about identifying and resolving other performance issues, see Performance.</p>
Determine the disk savings that are provided by data deduplication.	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Pools. 2. Select a pool and click Quick Look. 3. In the Data Deduplication area, view the Space saved row. 	<p>For advanced monitoring, to obtain detailed statistics about the data-deduplication process for a specific directory-container storage pool or cloud-container storage pool, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. Obtain a statistical report by issuing the GENERATE DEDUPSTATS command. Follow the instructions in GENERATE DEDUPSTATS (Generate data deduplication statistics for a directory-container storage pool). 3. View the statistical report by issuing the QUERY DEDUPSTATS command. Follow the instructions in QUERY DEDUPSTATS (Query data deduplication statistics).

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Verify that current backup files for device configuration and volume history information are saved.</p>	<p>Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations.</p> <p>To locate the volume history and device configuration files, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To locate the volume history and device configuration files, issue the following commands: <pre>query option volhistory query option devconfig</pre> 3. In the output, review the Option Setting column to find the file locations. <p>If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.</p>	

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Determine whether sufficient space is available for the instance directory file system.</p>	<p>Verify that at least 20% of free space is available in the instance directory file system. Take the action that is appropriate for your operating system:</p> <ul style="list-style-type: none"> • AIX To view available space in the file system, on the operating system command line, issue the following command: <pre>df -g instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> • Linux To view available space in the file system, on the operating system command line, issue the following command: <pre>df -h instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> • Windows In the Windows Explorer program, right-click the file system and click Properties. View the capacity information. <p>The preferred location of the instance directory depends on the operating system where the server is installed:</p> <ul style="list-style-type: none"> • AIX Linux /home/tsminst1/tsminst1 • Windows C:\tsminst1 <p>Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.</p>	
<p>Identify unexpected client activity.</p>	<p>To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Clients area. 2. To view activity over the past two weeks, double-click any client. 3. To view the number of bytes sent to the client, click the Properties tab. 4. In the Last Session area, view the Sent to client row. 	<p>When you double-click a client in the Clients table, the Activity over 2 Weeks area displays the amount of data that the client sent to the server each day.</p> <p>Periodically review the SQL activity summary table, which contains statistics about client sessions. To compare current activity with previous activity, use an SQL SELECT statement. If the level of activity is significantly different from previous activity, it might indicate a ransomware attack.</p> <p>Periodically review the activity log. Look for ANE messages that indicate how many files were backed up and inspected. Compare current data deduplication rates with previous rates. If an unusually high number of files were backed up, or the rate of data deduplication unexpectedly drops to 0, it might indicate a ransomware attack.</p>

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Monitor storage pool growth over time.</p>	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Pools area. 2. To view the capacity that was used over the last two weeks, select a pool and click Details. 	<p>Tips:</p> <ul style="list-style-type: none"> • To specify the time period that must elapse before all deduplicated extents are removed from a directory-container storage pool or cloud-container storage pool after they are no longer referenced by the inventory, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. 3. Specify the duration in the <code>Delay period for container reuse</code> field. • To determine data deduplication performance for directory-container and cloud-container storage pools, use the <code>GENERATE DEDUPSTATS</code> command. • To view data deduplication statistics for a storage pool, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. <p>Alternatively, use the <code>QUERY EXTENTUPDATES</code> command to display information about updates to data extents in directory-container or cloud-container storage pools. The command output can help you determine which data extents are no longer referenced and which ones are eligible to be deleted from the system. In the output, monitor the number of data extents that are eligible to be deleted from the system. This metric has a direct correlation to the amount of free space that will be available within the container storage pool.</p> • To display the amount of physical space that is occupied by a file space after the removal of the data deduplication savings, use the <code>select * from occupancy</code> command. The command output includes the <code>LOGICAL_MB</code> value. <code>LOGICAL_MB</code> is the amount of space that is used by the file space.
<p>Evaluate the timing of client schedules. Ensure that the start and end times of client schedules meet your business needs.</p>	<p>On the Operations Center Overview page, click Clients > Schedules.</p> <p>In the Schedules table, the Start column displays the configured start time for the scheduled operation. To see when the most recent operation was started, hover over the clock icon.</p>	<p>Tip: You can receive a warning message if a client operation runs longer than expected. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over Clients and click Schedules. 2. Select a schedule and click Details. 3. View the details of a schedule by clicking the blue arrow next to the row. 4. In the Run time alert field, specify the time when a warning message will be issued if the scheduled operation is not completed. 5. Click Save.

Task	Basic procedures	Advanced procedures and troubleshooting
Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks meet your business needs.	<p>On the Operations Center Overview page, click Servers > Maintenance.</p> <p>In the Maintenance table, review the information in the Last Run Time column. To see when the last maintenance task was started, hover over the clock icon.</p>	<p>Tip: If a maintenance task is running too long, change the start time or the maximum run time. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To change the start time or maximum run time for a task, issue the UPDATE SCHEDULE command. For instructions, see UPDATE SCHEDULE (Update a client schedule).

Related reference:

- [QUERY ACTLOG](#) (Query the activity log)
- [UPDATE STGPOOL](#) (Update a storage pool)
- [QUERY EXTENTUPDATES](#) (Query updated data extents)

Verifying license compliance

Verify that your IBM Spectrum Protect™ solution complies with the provisions of your licensing agreement. By verifying compliance regularly, you can track trends in data growth or processor value unit (PVU) usage. Use this information to plan for future license purchasing.

About this task

The method that you use to verify that your solution complies with license terms varies depending on the provisions of your IBM Spectrum Protect licensing agreement.

Front-end capacity licensing

The front-end model determines license requirements based on the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems.

Back-end capacity licensing

The back-end model determines license requirements based on the terabytes of data that are stored in primary storage pools and repositories.

Tips:

- To ensure the accuracy of front-end and back-end capacity estimates, install the most recent version of the client software on each client node.
- The front-end and back-end capacity information in the Operations Center is for planning and estimation purposes.

PVU licensing

The PVU model is based on the use of PVUs by server devices.

Important: The PVU calculations that are provided by IBM Spectrum Protect are considered estimates and are not legally binding. The PVU licensing information that is reported by IBM Spectrum Protect is not considered an acceptable substitute for the IBM® License Metric Tool.



For the most recent information about licensing models, see the information about product details and licenses at the IBM Spectrum Protect product family website. If you have questions or concerns about licensing requirements, contact your IBM Spectrum Protect software provider.

Procedure

To monitor license compliance, complete the steps that correspond to the provisions of your licensing agreement.

Tip: The Operations Center provides an email report that summarizes front-end and back-end capacity usage. Reports can be sent automatically to one or more recipients regularly. To configure and manage email reports, click Reports on the Operations Center menu bar.

Option	Description
--------	-------------

Option	Description
Front-end model	<p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>The front-end capacity estimate is displayed on the Front-end Usage page.</p> <p>b. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage.</p> <p>c. To estimate capacity for clients that did not report capacity usage, go to the following FTP site, which provides measuring tools and instructions:</p> <p><code>ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools</code></p> <p>To measure front-end capacity by script, complete the instructions in the most recently available licensing guide.</p> <p>d. Add the Operations Center estimate and any estimates that you obtained by using a script.</p> <p>e. Verify that the estimated capacity complies with your licensing agreement.</p>
Back-end model	<p>Restriction: If the source and target replication servers do not use the same policy settings, you cannot use the Operations Center to monitor back-end capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>b. Click the Back-end tab.</p> <p>c. Verify that the estimated amount of data complies with your licensing agreement.</p>
PVU model	<p>For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model.</p>

Tracking system status by using email reports

Set up the Operations Center to generate email reports that summarize system status. You can configure a mail server connection, change report settings, and optionally create custom reports.

Before you begin

Before you set up email reports, ensure that the following requirements are met:

- A Simple Mail Transfer Protocol (SMTP) host server is available to send and receive reports by email. The SMTP server must be configured as an open mail relay. You must also ensure that the IBM Spectrum Protect™ server that sends email messages has access to the SMTP server. If the Operations Center is installed on a separate computer, that computer does not require access to the SMTP server.
- To set up email reports, you must have system privilege for the server.
- To specify the recipients, you can enter one or more email addresses or administrator IDs. If you plan to enter an administrator ID, the ID must be registered on the hub server and must have an email address that is associated with it. To specify an email address for an administrator, use the EMAILADDRESS parameter of the UPDATE ADMIN command.

About this task

You can configure the Operations Center to send a general operations report, a license compliance report, and one or more custom reports. You create custom reports by selecting a template from a set of commonly used report templates or by entering SQL SELECT statements to query managed servers.

Procedure

To set up and manage email reports, complete the following steps:

1. On the Operations Center menu bar, click Reports.
2. If an email server connection is not yet configured, click Configure Mail Server and complete the fields. After you configure the mail server, the general operations report and license compliance report are enabled.

3. To change report settings, select a report, click Details, and update the form.
4. Optional: To add a custom report, click + Report, and complete the fields.
Tip: To immediately run and send a report, select the report and click Send.

Results

Enabled reports are sent according to the specified settings.

Related reference:

[UPDATE ADMIN](#) (Update an administrator)

Managing operations for a single-site disk solution

Use this information to manage operations for a single-site disk solution with IBM Spectrum Protect™ that includes a server and uses data deduplication for a single location.

- **Managing the Operations Center**
The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.
- **Protecting applications, virtual machines, and systems**
The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.
- **Managing data storage**
Manage your data for efficiency and add supported devices and media to the server to store client data.
- **Securing the IBM Spectrum Protect server**
Secure the IBM Spectrum Protect server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.
- **Stopping and starting the server**
Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.
- **Planning to upgrade the server**
When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.
- **Preparing for an outage or system update**
Prepare IBM Spectrum Protect to maintain your system in a consistent state during a planned power outage or system update.
- **Implementing a disaster recovery plan**
Implement a disaster recovery strategy to recover your applications if a disaster occurs and to ensure high server availability.
- **Recovering from system outages**
For IBM Spectrum Protect single-site disk solutions, you can recover the inventory locally only and restore the database to protect your data.

Managing the Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect™ environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

- **Adding and removing spoke servers**
In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.
- **Starting and stopping the web server**
The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.
- **Restarting the initial configuration wizard**
You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.
- **Changing the hub server**
You can use the Operations Center to remove the hub server of IBM Spectrum Protect, and configure another hub server.

- Restoring the configuration to the preconfiguration state
If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect servers are not defined as hub or spoke servers.

Adding and removing spoke servers

In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

About this task

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

- Adding a spoke server
After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.
- Removing a spoke server
You can remove a spoke server from the Operations Center.

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of the hub server.

Procedure

1. In the Operations Center menu bar, click Servers. The Servers page opens.

In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke server.

2. Complete one of the following steps:
 - Click the server to highlight it, and in the table menu bar, click Monitor Spoke.
 - If the server that you want to add is not shown in the table, and secure SSL/TLS communication is not required, click + Spoke in the table menu bar.
3. Provide the necessary information, and complete the steps in the spoke configuration wizard.
Tip: If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Removing a spoke server

You can remove a spoke server from the Operations Center.

About this task

You might need to remove a spoke server in the following situations, for example:

- You want to move the spoke server from one hub server to another hub server.
- You want to decommission the spoke server.

Procedure

To remove the spoke server from the group of servers that are managed by the hub server, complete the following steps:

1. From the IBM Spectrum Protect™ command line, issue the following command on the hub server:

```
QUERY MONITORSETTINGS
```

2. From the output of the command, copy the name that is in the Monitored Group field.
3. Issue the following command on the hub server, where *group_name* represents the name of the monitored group, and *member_name* represents the name of the spoke server:

```
DELETE GRPMEMBER group_name member_name
```

4. Optional: If you want to move the spoke server from one hub server to another hub server, do **not** complete this step. Otherwise, you can disable alerting and monitoring on the spoke server by issuing the following commands on the spoke server:

```
SET STATUSMONITOR OFF  
SET ALERTMONITOR OFF
```

5. Optional: If the spoke server definition is used for other purposes, such as enterprise configuration, command routing, storing virtual volumes, or library management, do **not** complete this step. Otherwise, you can delete the spoke server definition on the hub server by issuing the following command on the hub server:

```
DELETE SERVER spoke_server_name
```

Tip: If a server definition is deleted immediately after the server is removed from the monitored group, status information for the server can remain in the Operations Center indefinitely.

To avoid this issue, wait until the status collection interval passes before you delete the server definition. The status collection interval is shown on the Settings page of the Operations Center.

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.

Procedure

1. Stop the web server.
 - o **AIX** From the */installation_dir/ui/utls* directory, where *installation_dir* represents the directory where the Operations Center is installed, issue the following command:

```
./stopserver.sh
```
 - o **Linux** Issue the following command:

```
service opscenter.rc stop
```
 - o **Windows** From the Services window, stop the IBM Spectrum Protect™ Operations Center service.
2. Start the web server.
 - o **AIX** From the */installation_dir/ui/utls* directory, where *installation_dir* represents the directory where the Operations Center is installed, issue the following command:

```
./startserver.sh
```
 - o **Linux** Issue the following commands:

Start the server:

```
service opscenter.rc start
```


Restart the server:

```
service opscenter.rc restart
```


Determine whether the server is running:

```
service opscenter.rc status
```
 - o **Windows** From the Services window, start the IBM Spectrum Protect Operations Center service.

Restarting the initial configuration wizard

You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.

Before you begin

To change the following settings, use the Settings page in the Operations Center rather than restarting the initial configuration wizard:

- The frequency at which status data is refreshed
- The duration that alerts remain active, inactive, or closed
- The conditions that indicate that clients are at risk

The Operations Center help includes more information about how to change these settings.

About this task

To restart the initial configuration wizard, you must delete a properties file that includes information about the hub server connection. However, any alerting, monitoring, at-risk, or multiserver settings that were configured for the hub server are not deleted. These settings are used as the default settings in the configuration wizard when the wizard restarts.

Procedure

1. Stop the Operations Center web server.
 2. On the computer where the Operations Center is installed, go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:
 - o **AIX** | **Linux** *installation_dir*/ui/Liberty/usr/servers/guiServer
 - o **Windows** *installation_dir*\ui\Liberty\usr\servers\guiServer
- For example:
- o **AIX** | **Linux** /opt/tivoli/tsm/ui/Liberty/usr/servers/guiServer
 - o **Windows** c:\Program Files\Tivoli\TSM\ui\Liberty\usr\servers\guiServer
3. In the guiServer directory, delete the serverConnection.properties file.
 4. Start the Operations Center web server.
 5. Open the Operations Center.
 6. Use the configuration wizard to reconfigure the Operations Center. Specify a new password for the monitoring administrator ID.
 7. On any spoke servers that were previously connected to the hub server, update the password for the monitoring administrator ID by issuing the following command from the IBM Spectrum Protect™ command-line interface:

```
UPDATE ADMIN IBM-OC-hub_server_name new_password
```

Restriction: Do not change any other settings for this administrator ID. After you specify the initial password, this password is managed automatically by the Operations Center.

Changing the hub server

You can use the Operations Center to remove the hub server of IBM Spectrum Protect™, and configure another hub server.

Procedure

1. Restart the initial configuration wizard of the Operations Center. As part of this procedure, you delete the existing hub server connection.
2. Use the wizard to configure the Operations Center to connect to the new hub server.

Related tasks:

Restarting the initial configuration wizard

Restoring the configuration to the preconfiguration state

If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect™ servers are not defined as hub or spoke servers.

Procedure

To restore the configuration, complete the following steps:

1. Stop the Operations Center web server.
2. Unconfigure the hub server by completing the following steps:
 - a. On the hub server, issue the following commands:

```
SET MONITORINGADMIN ""
SET MONITOREDSEVERGROUP ""
SET STATUSMONITOR OFF
SET ALERTMONITOR OFF
REMOVE ADMIN IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. Reset the password for the hub server by issuing the following command on the hub server:

```
SET SERVERPASSWORD ""
```

Attention: Do not complete this step if the hub server is configured with other servers for other purposes, such as library sharing, exporting and importing of data, or node replication.

3. Unconfigure any spoke servers by completing the following steps:
 - a. On the hub server, to determine whether any spoke servers remain as members of the server group, issue the following command:

```
QUERY SERVERGROUP IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the name of the monitored server group that was automatically created when you configured the first spoke server. This server group name is also the same as the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. On the hub server, to delete spoke servers from the server group, issue the following command for each spoke server:

```
DELETE GRPMEMBER IBM-OC-hub_server_name spoke_server_name
```

- c. After all spoke servers are deleted from the server group, issue the following commands on the hub server:

```
DELETE SERVERGROUP IBM-OC-hub_server_name
SET MONITOREDSEVERGROUP ""
```

- d. On each spoke server, issue the following commands:

```
REMOVE ADMIN IBM-OC-hub_server_name
SETOPT PUSHSTATUS NO
SET ALERTMONITOR OFF
SET STATUSMONITOR OFF
```

- e. On each spoke server, delete the definition of the hub server by issuing the following command:

```
DELETE SERVER hub_server_name
```

Attention: Do not complete this step if the definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

- f. On the hub server, delete the definition of each spoke server by issuing the following command:

```
DELETE SERVER spoke_server_name
```

Attention: Do not complete this step if the server definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

4. Restore the default settings on each server by issuing the following commands:

```
SET STATUSREFRESHINTERVAL 5
SET ALERTUPDATEINTERVAL 10
SET ALERTACTIVEDURATION 480
SET ALERTINACTIVEDURATION 480
SET ALERTCLOSEDDURATION 60
SET STATUSATRISKINTERVAL TYPE=AP INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=VM INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=SY INTERVAL=24
SET STATUSSKIPASFAILURE YES TYPE=ALL
```

5. Restart the initial configuration wizard of the Operations Center.

Related tasks:

Restarting the initial configuration wizard

Starting and stopping the web server

Protecting applications, virtual machines, and systems

The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.

- **Adding clients**
After you implement a data protection solution with IBM Spectrum Protect, you can expand the solution by adding clients.
- **Managing client operations**
You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.
- **Managing client upgrades**
When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.
- **Decommissioning a client node**
If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect server, but the workstation is no longer used, you can decommission the workstation.
- **Deactivating data to free storage space**
In some cases, you can deactivate data that is stored on the IBM Spectrum Protect server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

Adding clients

After you implement a data protection solution with IBM Spectrum Protect™, you can expand the solution by adding clients.

About this task

The procedure describes basic steps for adding a client. For more specific instructions about configuring clients, see the documentation for the product that you install on the client node. You can have the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

Procedure

To add a client, complete the following steps:

1. Select the software to install on the client node and plan the installation. Follow the instructions in [Selecting the client software and planning the installation](#).
2. Specify how to back up and archive client data. Follow the instructions in [Specifying rules for backing up and archiving client data](#).
3. Specify when to back up and archive client data. Follow the instructions in [Scheduling backup and archive operations](#).

4. To allow the client to connect to the server, register the client. Follow the instructions in Registering clients.
5. To start protecting a client node, install and configure the selected software on the client node. Follow the instructions in Installing and configuring clients.

Selecting the client software and planning the installation

Different types of data require different types of protection. Identify the type of data that you must protect and select the appropriate software.

About this task

The preferred practice is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you install a product for which the client acceptor does not run schedules, you must follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

Procedure

Based on your goal, select the products to install and review the installation instructions.

Tip: If you install the client software now, you must also complete the client configuration tasks that are described in Installing and configuring clients before you can use the client.

Goal	Product and description	Installation instructions
Protect a file server or workstation	The backup-archive client backs up and archives files and directories from file servers and workstations to storage. You can also restore and retrieve backup versions and archived copies of files.	<ul style="list-style-type: none"> • Backup-archive client requirements • Install UNIX and Linux backup-archive clients • Installing the Windows client for the first time
Protect applications with snapshot backup and restore capabilities	IBM Spectrum Protect Snapshot protects data with integrated, application-aware snapshot backup and restore capabilities. You can protect data that is stored by IBM DB2® database software and SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications.	<ul style="list-style-type: none"> • Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux • Installing and upgrading IBM Spectrum Protect Snapshot for VMware • Installing and upgrading IBM Spectrum Protect Snapshot for Windows
Protect an email application on an IBM Domino® server	IBM Spectrum Protect for Mail: Data Protection for IBM® Domino automates data protection so that backups are completed without shutting down IBM Domino servers.	<ul style="list-style-type: none"> • Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) • Installation of Data Protection for IBM Domino on a Windows system (V7.1.0)
Protect an email application on a Microsoft Exchange server	IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server automates data protection so that backups are completed without shutting down Microsoft Exchange servers.	Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
Protect an IBM DB2 database	The application programming interface (API) of the backup-archive client can be used to back up DB2 data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)
Protect an IBM Informix® database	The API of the backup-archive client can be used to back up Informix data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)

Goal	Product and description	Installation instructions
Protect a Microsoft SQL database	IBM Spectrum Protect for Databases: Data Protection for Microsoft SQL Server protects Microsoft SQL data.	Installing Data Protection for SQL Server on Windows Server Core
Protect an Oracle database	IBM Spectrum Protect for Databases: Data Protection for Oracle protects Oracle data.	Data Protection for Oracle installation
Protect an SAP environment	IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP provides protection that is customized for SAP environments. The product is designed to improve the availability of SAP database servers and reduce administration workload.	<ul style="list-style-type: none"> Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2 Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
Protect a virtual machine	<p>IBM Spectrum Protect for Virtual Environments provides protection that is tailored for Microsoft Hyper-V and VMware virtual environments. You can use IBM Spectrum Protect for Virtual Environments to create incremental forever backups that are stored on a centralized server, create backup policies, and restore virtual machines or individual files.</p> <p>Alternatively, use the backup-archive client to back up and restore a full VMware or Microsoft Hyper-V virtual machine. You can also back up and restore files or directories from a VMware virtual machine.</p>	<ul style="list-style-type: none"> Installing Data Protection for Microsoft Hyper-V Installing and upgrading Data Protection for VMware Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)

Tip: To use the client for space management, you can install IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows.

Specifying rules for backing up and archiving client data

Before you add a client, ensure that appropriate rules are specified for backup and archive operations for the client data. During the client registration process, you assign the client node to a policy domain, which has the rules that control how and when client data is stored.

Before you begin

Determine how to proceed:

- If you are familiar with the policies that are configured for your solution and you know that they do not require changes, continue with Scheduling backup and archive operations.
- If you are not familiar with the policies, follow the steps in this procedure.

About this task

Policies affect how much data is stored over time, and how long data is retained and available for clients to restore. To meet objectives for data protection, you can update the default policy and create your own policies. A policy includes the following rules:

- How and when files are backed up and archived to server storage.
- The number of copies of a file and the length of time copies are kept in server storage.

During the client registration process, you assign a client to a *policy domain*. The policy for a specific client is determined by the rules in the policy domain to which the client is assigned. In the policy domain, the rules that are in effect are in the active *policy set*.

When a client backs up or archives a file, the file is bound to a management class in the active policy set of the policy domain. A *management class* is the key set of rules for managing client data. The backup and archive operations on the client use the

settings in the default management class of the policy domain unless you further customize policy. A policy can be customized by defining more management classes and assigning their use through client options.

Client options can be specified in a local, editable file on the client system and in a client option set on the server. The options in the client option set on the server can override or add to the options in the local client option file.

Procedure

1. Review the policies that are configured for your solution by following the instructions in Viewing policies.
2. If you need to make minor changes to meet data retention requirements, follow the instructions in Editing policies.
3. Optional: If you need to create policy domains or make extensive changes to policies to meet data retention requirements, see Customizing policies.

Viewing policies

View policies to determine whether they must be edited to meet your requirements.

Procedure

1. To view the active policy set for a policy domain, complete the following steps:
 - a. On the Services page of the Operations Center, select a policy domain and click Details.
 - b. On the Summary page for the policy domain, click the Policy Sets tab.

Tip: To help ensure that you can recover data after a ransomware attack, apply the following guidelines:

 - Ensure that the value in the Backups column is a minimum of 2. The preferred value is 3, 4, or more.
 - Ensure that the value in the Keep Extra Backups column is a minimum of 14 days. The preferred value is 30 or more days.
 - Ensure that the value in the Keep Archives column is a minimum of 30 days.

If IBM Spectrum Protect™ for Space Management software is installed on the client, ensure that data is backed up before you migrate it. On the DEFINE MGMTCLASS or UPDATE MGMTCLASS command, specify MIGREQUIRESBKUP=YES. Then, follow the guidelines in the tip.
2. To view inactive policy sets for a policy domain, complete the following steps:
 - a. On the Policy Sets page, click the Configure toggle. You can now view and edit the policy sets that are inactive.
 - b. Scroll through the inactive policy sets by using the forward and back arrows. When you view an inactive policy set, the settings that differentiate the inactive policy set from the active policy set are highlighted.
 - c. Click the Configure toggle. The policy sets are no longer editable.

Editing policies

To change the rules that apply to a policy domain, edit the active policy set for the policy domain. You can also activate a different policy set for a domain.

Before you begin

Changes to policy can affect data retention. Ensure that you continue to back up data that is essential to your organization so that you can restore that data if a disaster occurs. Also, ensure that your system has sufficient storage space for planned backup operations.

About this task

You edit a policy set by changing one or more management classes within the policy set. If you edit the active policy set, the changes are not available to clients unless you reactivate the policy set. To make the edited policy set available to clients, activate the policy set.

Although you can define multiple policy sets for a policy domain, only one policy set can be active. When you activate a different policy set, it replaces the currently active policy set.

To learn about preferred practices for defining policies, see Customizing policies.

Procedure

1. On the Services page of the Operations Center, select a policy domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab.

The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.

3. Click the Configure toggle. The policy set is editable.
4. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.
5. Edit the policy set by completing any of the following actions:

Option	Description
Add a management class	<ol style="list-style-type: none"> a. In the Policy Sets table, click +Management Class. b. To specify the rules for backing up and archiving data, complete the fields in the Add Management Class window. c. To make the management class the default management class, select the Make default check box. d. Click Add.
Delete a management class	<p>In the Management Class column, click -.</p> <p>Tip: To delete the default management class, you must first assign a different management class as the default.</p>
Make a management class the default management class	<p>In the Default column for the management class, click the radio button.</p> <p>Tip: The default management class manages client files when another management class is not assigned to, or appropriate for managing, a file. To ensure that clients can always back up and archive files, choose a default management class that contains rules for both backing up and archiving files.</p>
Modify a management class	To change the properties of a management class, update the fields in the table.

6. Click Save.

Attention: When you activate a new policy set, data might be lost. Data that is protected under one policy set might not be protected under another policy set. Therefore, before you activate a policy set, ensure that the differences between the previous policy set and the new policy set do not cause data to be lost.
7. Click Activate. A summary of the differences between the active policy set and the new policy set is displayed. Ensure that the changes in the new policy set are consistent with your data retention requirements by completing the following steps:
 - a. Review the differences between corresponding management classes in the two policy sets, and consider the consequences for client files. Client files that are bound to management classes in the active policy set will be bound to the management classes with the same names in the new policy set.
 - b. Identify management classes in the active policy set that do not have counterparts in the new policy set, and consider the consequences for client files. Client files that are bound to these management classes will be managed by the default management class in the new policy set.
 - c. If the changes to be implemented by the policy set are acceptable, select the I understand that these updates can cause data loss check box and click Activate.

Scheduling backup and archive operations

Before you register a new client with the server, ensure that a schedule is available to specify when backup and archive operations take place. During the registration process, you assign a schedule to the client.

Before you begin

Determine how to proceed:

- If you are familiar with the schedules that are configured for the solution and you know that they do not require modification, continue with Registering clients.
- If you are not familiar with the schedules or the schedules require modification, follow the steps in this procedure.


About this task

Typically, backup operations for all clients must be completed daily. Schedule client and server workloads to achieve the best performance for your storage environment. To avoid the overlap of client and server operations, consider scheduling client backup and archive operations so that they run at night. If client and server operations overlap or are not given enough time and resources to be processed, you might experience decreased system performance, failed operations, and other issues.

Procedure

1. Review available schedules by hovering over Clients on the Operations Center menu bar. Click Schedules.
2. Optional: Modify or create a schedule by completing the following steps:

Option	Description
Modify a schedule	<ol style="list-style-type: none">a. In the Schedules view, select the schedule and click Details.b. On the Schedule Details page, view details by clicking the blue arrows at the beginning of the rows.c. Modify the settings in the schedule, and click Save.
Create a schedule	In the Schedules view, click +Schedule and complete the steps to create a schedule.

3. Optional: To configure schedule settings that are not visible in the Operations Center, use a server command. For example, you might want to schedule a client operation that backs up a specific directory and assigns it to a management class other than the default.
 - a. On the Overview page of the Operations Center, hover over the settings icon  and click Command Builder.
 - b. Issue the DEFINE SCHEDULE command to create a schedule or the UPDATE SCHEDULE command to modify a schedule. For more information about the commands, see DEFINE SCHEDULE (Define a schedule for an administrative command) or UPDATE SCHEDULE (Update a client schedule).

Related tasks:

[🔗 Tuning the schedule for daily operations](#)

Registering clients

Register a client to ensure that the client can connect to the server, and the server can protect client data.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see technote 7048963.

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see Authenticating users by using an Active Directory database.

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the REGISTER NODE command and specify the USERID parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see Register a node.

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:
 - a. On the Operations Center menu bar, click Clients.
 - b. In the Clients table, click +Client.
 - c. Complete the steps in the Add Client wizard:
 - i. Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the Enable check box.
 - ii. In the Configuration window, copy the TCPSERVERADDRESS, TCPPORT, NODENAME, and DEDUPLICATION option values.
Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
 - iii. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
 - iv. Set how risks are displayed for the client by specifying the at-risk setting.
 - v. Click Add Client.

Related reference:

- 🔗 [Tcpserveraddress option](#)
- 🔗 [Tcpsport option](#)
- 🔗 [Nodename option](#)
- 🔗 [Deduplication option](#)

Installing and configuring clients

To start protecting a client node, you must install and configure the selected software.

Procedure

If you already installed the software, start at step 2.

1. Take one of the following actions:

- To install software on an application or client node, follow the instructions.

Software	Link to instructions
IBM Spectrum Protect™ backup-archive client	<ul style="list-style-type: none"> ▪ Install UNIX and Linux backup-archive clients ▪ Installing the Windows client for the first time <p>Tip: You can also update existing clients by using the Operations Center. For instructions, see Scheduling client updates.</p>
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> ▪ Data Protection for Oracle installation ▪ Installing Data Protection for SQL Server on Windows Server Core
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> ▪ Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) ▪ Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) ▪ Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> ▪ Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux ▪ Installing and upgrading IBM Spectrum Protect Snapshot for VMware ▪ Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle

- To install software on a virtual machine client node, follow the instructions for the selected backup type.

Backup type	Link to instructions
If you plan to create full VMware backups of virtual machines, install and configure the IBM Spectrum Protect backup-archive client.	<ul style="list-style-type: none"> ▪ Install UNIX and Linux backup-archive clients ▪ Installing the Windows client for the first time
If you plan to create incremental forever full backups of virtual machines, install and configure IBM Spectrum Protect for Virtual Environments and the backup-archive client on the same client node or on different client nodes.	<ul style="list-style-type: none"> ▪ IBM Spectrum Protect for Virtual Environments online product documentation <p>Tip: You can obtain the software for IBM Spectrum Protect for Virtual Environments and the backup-archive client in the IBM Spectrum Protect for Virtual Environments installation package.</p>

2. To allow the client to connect to the server, add or update the values for the TCPSERVERADDRESS, TCPPORT, and NODENAME options in the client options file. Use the values that you recorded when you registered the client (Registering clients).
- For clients that are installed on an AIX®, Linux, or Mac OS X operating system, add the values to the client system-options file, dsm.sys.
 - For clients that are installed on a Windows operating system, add the values to the dsm.opt file.

- By default, the options files are in the installation directory.
3. If you installed a backup-archive client on a Linux or Windows operating system, install the client management service on the client. Follow the instructions in [Installing the client management service](#).
 4. Configure the client to run scheduled operations. Follow the instructions in [Configuring the client to run scheduled operations](#).
 5. Optional: Configure communications through a firewall. Follow the instructions in [Configuring client/server communications through a firewall](#).
 6. Run a test backup to verify that data is protected as you planned. For example, for a backup-archive client, complete the following steps:
 - a. On the Clients page of the Operations Center, select the client that you want to back up, and click Back Up.
 - b. Verify that the backup completes successfully and that there are no warning or error messages.
 7. Monitor the results of the scheduled operations for the client in the Operations Center.

What to do next

If you need to change what is getting backed up from the client, follow the instructions in [Modifying the scope of a client backup](#).

Configuring the client to run scheduled operations

You must configure and start a client scheduler on the client node. The client scheduler enables communication between the client and server so that scheduled operations can occur. For example, scheduled operations typically include backing up files from a client.

About this task

The preferred method is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations. The client acceptor manages the client scheduler so that the scheduler runs only when required:

- When it is time to query the server about the next scheduled operation
- When it is time to start the next scheduled operation

By using the client acceptor, you can reduce the number of background processes on the client and help to avoid memory retention problems.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you installed a product for which the client acceptor does not run schedules, follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

If your business uses a third-party scheduling tool as standard practice, you can use that scheduling tool as an alternative to the client acceptor. Typically, third-party scheduling tools start client programs directly by using operating system commands. To configure a third-party scheduling tool, see the product documentation.

Procedure

To configure and start the client scheduler by using the client acceptor, follow the instructions for the operating system that is installed on the client node:

AIX® and Oracle Solaris

- a. From the backup-archive client GUI, click Edit > Client Preferences.
- b. Click the Web Client tab.
- c. In the Managed Services Options field, click Schedule. If you also want the client acceptor to manage the web client, click the Both option.
- d. To ensure that the scheduler can start unattended, in the `dsm.sys` file, set the `passwordaccess` option to `generate`.
- e. To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

- f. Start the client acceptor by issuing the following command on the command line:


```
/usr/bin/dsmcad
```

- g. To enable the client acceptor to start automatically after a system restart, add the following entry to the system startup file (typically, `/etc/inittab`):

```
tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # Client Acceptor Daemon
```

Linux

- From the backup-archive client GUI, click Edit > Client Preferences.
- Click the Web Client tab.
- In the Managed Services Options field, click Schedule. If you also want the client acceptor to manage the web client, click the Both option.
- To ensure that the scheduler can start unattended, in the `dsm.sys` file, set the `passwordaccess` option to `generate`.
- To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

- f. Start the client acceptor by logging in with the root user ID and issuing the following command:

```
service dsmcad start
```

- g. To enable the client acceptor to start automatically after a system restart, add the service by issuing the following command at a shell prompt:

```
# chkconfig --add dsmcad
```

MAC OS X

- In the backup-archive client GUI, click Edit > Client Preferences.
- To ensure that the scheduler can start unattended, click Authorization, select Password Generate, and click Apply.
- To specify how services are managed, click Web Client, select Schedule, click Apply, and click OK.
- To ensure that the generated password is saved, restart the backup-archive client.
- Use the IBM Spectrum Protect Tools for Administrators application to start the client acceptor.

Windows

- In the backup-archive client GUI, click Utilities > Setup Wizard > Help me configure the Client Scheduler. Click Next.
- Read the information on the Scheduler Wizard page and click Next.
- On the Scheduler Task page, select Install a new or additional scheduler and click Next.
- On the Scheduler Name and Location page, specify a name for the client scheduler that you are adding. Then, select Use the Client Acceptor daemon (CAD) to manage the scheduler and click Next.
- Enter the name that you want to assign to this client acceptor. The default name is Client Acceptor. Click Next.
- Complete the configuration by stepping through the wizard.
- Update the client options file, `dsm.opt`, and set the `passwordaccess` option to `generate`.
- To store the client node password, issue the following command at the command prompt:

```
dsmc query sess
```

Enter the client node password when prompted.

- Start the client acceptor service from the Services Control page. For example, if you used the default name, start the Client Acceptor service. Do not start the scheduler service that you specified on the Scheduler Name and Location page. The scheduler service is started and stopped automatically by the client acceptor service as needed.

Configuring client/server communications through a firewall

If a client must communicate with a server through a firewall, you must enable client/server communications through the firewall.

Before you begin

If you used the Add Client wizard to register a client, find the option values in the client options file that you obtained during that process. You can use the values to specify ports.

About this task

Attention: Do not configure a firewall in a way that might cause termination of sessions that are in use by a server or storage agent. Termination of a valid session can cause unpredictable results. Processes and sessions might appear to stop due to input/output errors. To help exclude sessions from timeout restrictions, configure known ports for IBM Spectrum Protect™ components. Ensure that the KEEPALIVE server option remains set to the default value of YES. In this way, you can help to ensure that client/server communication is uninterrupted. For instructions about setting the KEEPALIVE server option, see KEEPALIVE.

Procedure

Open the following ports to allow access through the firewall:

TCP/IP port for the backup-archive client, command-line administrative client, and the client scheduler

Specify the port by using the tcpport option in the client options file. The tcpport option in the client options file must match the TCPSPORT option in the server options file. The default value is 1500. If you decide to use a value other than the default, specify a number in the range 1024 - 32767.

HTTP port to enable communication between the web client and remote workstations

Specify the port for the remote workstation by setting the httpport option in the client options file of the remote workstation. The default value is 1581.

TCP/IP ports for the remote workstation

The default value of 0 (zero) causes two free port numbers to be randomly assigned to the remote workstation. If you do not want the port numbers to be randomly assigned, specify values by setting the webports option in the client options file of the remote workstation.

TCP/IP port for administrative sessions

Specify the port on which the server waits for requests for administrative client sessions. The value of the client tcpadminport option must match the value of the TCPADMINPORT server option. In this way, you can secure administrative sessions within a private network.

Managing client operations

You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.

About this task

In some cases, you can resolve client errors by stopping and starting the client acceptor. If client nodes or administrator IDs are locked, you can resolve the issue by unlocking the client node or administrator ID, and then resetting the password.

For detailed instructions about identifying and resolving client errors, see Resolving client problems.

- Evaluating errors in client error logs
You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.
- Stopping and restarting the client acceptor
If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.
- Resetting passwords
If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.
- Modifying the scope of a client backup
When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

Evaluating errors in client error logs

You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

Before you begin

To resolve errors in a backup-archive client on a Linux or Windows operating system, ensure that the client management service is installed and started. For installation instructions, see [Installing the client management service](#). For instructions about verifying the installation, see [Verifying that the client management service is installed correctly](#).

Procedure

To diagnose and resolve client errors, take one of the following actions:

- If the client management service is installed on the client node, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click Details.
 3. On the client Summary page, click the Diagnosis tab.
 4. Review the retrieved log messages.

Tips:

- To show or hide the Client Logs pane, double-click the Client Logs bar.
- To resize the Client Logs pane, click and drag the Client Logs bar.

If suggestions are displayed on the Diagnosis page, select a suggestion. In the Client Logs pane, client log messages to which the suggestion relates are highlighted.

5. Use the suggestions to resolve the problems that are indicated by the error messages.
Tip: Suggestions are provided for only a subset of client messages.
- If the client management service is not installed on the client node, review the error logs for the installed client.

Stopping and restarting the client acceptor

If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.

About this task

In some cases, you can resolve client scheduling problems by stopping and restarting the client acceptor. The client acceptor must be running to ensure that scheduled operations can occur on the client. For example, if you change the IP address or domain name of the server, you must restart the client acceptor.

Procedure

Follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

- To stop the client acceptor, complete the following steps:
 - a. Determine the process ID for the client acceptor by issuing the following command on the command line:

```
ps -ef | grep dsmcad
```

Review the output. In the following sample output, 6764 is the process ID for the client acceptor:

```
root 6764 1 0 16:26:35 ? 0:00 /usr/bin/dsmcad
```

- b. Issue the following command on the command line:

```
kill -9 PID
```

where *PID* specifies the process ID for the client acceptor.

- To start the client acceptor, issue the following command on the command line:

```
/usr/bin/dsmcad
```

Linux

- To stop the client acceptor (and not restart it), issue the following command:

```
# service dsmcad stop
```

- To stop and restart the client acceptor, issue the following command:

```
# service dsmscad restart
```

MAC OS X

Click Applications > Utilities > Terminal.

- To stop the client acceptor, issue the following command:

```
/bin/launchctl unload -w com.ibm.tivoli.dsmscad
```

- To start the client acceptor, issue the following command:

```
/bin/launchctl load -w com.ibm.tivoli.dsmscad
```

Windows

- To stop the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Stop and OK.
- To restart the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Start and OK.

Related reference:

[Resolving client scheduling problems](#)

Resetting passwords

If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

Procedure

To resolve password issues, take one of the following actions:

- If a backup-archive client is installed on a client node, and the password is lost or forgotten, complete the following steps:
 1. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwnreset=yes
```

where *node_name* specifies the client node and *new_password* specifies the password that you assign.

2. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.
Tip: The password is generated automatically if you previously set the passwordaccess option to *generate* in the client options file.

- If an administrator is locked out because of password issues, complete the following steps:
 1. To provide the administrator with access to the server, issue the UNLOCK ADMIN command. For instructions, see UNLOCK ADMIN (Unlock an administrator).
 2. Set a new password by using the UPDATE ADMIN command:

```
update admin admin_name new_password forcepwnreset=yes
```

where *admin_name* specifies the name of the administrator and *new_password* specifies the password that you assign.

- If a client node is locked, complete the following steps:
 1. Determine why the client node is locked and whether it must be unlocked. For example, if the client node is decommissioned, the client node is being removed from the production environment. You cannot reverse the decommission operation, and the client node remains locked. A client node also might be locked if the client data is the subject of a legal investigation.
 2. If you must unlock a client node, use the UNLOCK NODE command. For instructions, see UNLOCK NODE (Unlock a client node).

3. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwreset=yes
```

where *node_name* specifies the name of the node and *new_password* specifies the password that you assign.

4. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to `generate` in the client options file.

Modifying the scope of a client backup

When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

About this task

When you exclude unnecessary objects from backup operations, you get better control of the amount of storage space that is required for backup operations, and the cost of storage. Depending on your licensing package, you also might be able to limit licensing costs.

Procedure

How you modify the scope of backup operations depends on the product that is installed on the client node:

- For a backup-archive client, you can create an include-exclude list to include or exclude a file, groups of files, or directories from backup operations. To create an include-exclude list, follow the instructions in [Creating an include-exclude list](#).

To ensure consistent use of an include-exclude list for all clients of one type, you can create a client option set on the server that contains the required options. Then, you assign the client option set to each of the clients of the same type. For details, see [Controlling client operations through client option sets](#).

- For a backup-archive client, you can specify the objects to include in an incremental backup operation by using the domain option. Follow the instructions in [Domain option](#).
- For other products, to define which objects are included in and excluded from backup operations, follow the instructions in the product documentation.

Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Before you begin

1. Review the client/server compatibility requirements in [technote 1053218](#). If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted.
2. Verify system requirements for the client in [IBM Spectrum Protect™ Supported Operating Systems](#).
3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with servers that are configured as library managers. See [technote 1302789](#).

If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, complete the instructions that are listed in the following table.

Software	Link to instructions
IBM Spectrum Protect backup-archive client	<ul style="list-style-type: none">• Scheduling client updates

Software	Link to instructions
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux Installing and upgrading IBM Spectrum Protect Snapshot for VMware Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> Upgrading Data Protection for SQL Server Data Protection for Oracle installation Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Virtual Environments	<ul style="list-style-type: none"> Installing and upgrading Data Protection for VMware Installing Data Protection for Microsoft Hyper-V

Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect™ server, but the workstation is no longer used, you can decommission the workstation.

About this task

When you start the decommission process, the server locks the client node to prevent it from accessing the server. Files that belong to the client node are gradually deleted, and then the client node is deleted. You can decommission the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

The simplest method for decommissioning a client node is to use the Operations Center. The decommission process runs in the background. If the client is configured to replicate client data, the Operations Center automatically removes the client from replication on the source and target replication servers before it decommissions the client.

Tip: Alternatively, you can decommission a client node by issuing the `DECOMMISSION NODE` or `DECOMMISSION VM` command. You might want to use this method in the following cases:

- To schedule the decommission process for the future or to run a series of commands by using a script, specify the decommission process to run in the background.
- To monitor the decommission process for debugging purposes, specify the decommission process to run in the foreground. If you run the process in the foreground, you must wait for the process to be completed before you continue with other tasks.

Procedure

Take one of the following actions:

- To decommission a client in the background by using the Operations Center, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click More > Decommission.
- To decommission a client node by using an administrative command, take one of the following actions:
 - To decommission an application or system client node in the background, issue the DECOMMISSION NODE command. For example, if the client node is named AUSTIN, issue the following command:


```
decommission node austin
```
 - To decommission an application or system client node in the foreground, issue the DECOMMISSION NODE command and specify the `wait=yes` parameter. For example, if the client node is named AUSTIN, issue the following command:


```
decommission node austin wait=yes
```
 - To decommission a virtual machine in the background, issue the DECOMMISSION VM command. For example, if the virtual machine is named AUSTIN, the file space is 7, and the file space name is specified by the file space ID, issue the following command:


```
decommission vm austin 7 nametype=fsid
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid
```
 - To decommission a virtual machine in the foreground, issue the DECOMMISSION VM command and specify the `wait=yes` parameter. For example, issue the following command:


```
decommission vm austin 7 nametype=fsid wait=yes
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid wait=yes
```

What to do next

Watch for error messages, which might be displayed in the user interface or in the command output, immediately after you run the process.

You can verify that the client node is decommissioned:

1. On the Operations Center Overview page, click Clients.
2. In the Clients table, in the At risk column, review the state:
 - A DECOMMISSIONED state specifies that the node is decommissioned.
 - A null value specifies that the node is not decommissioned.
 - A PENDING state specifies that the node is being decommissioned, or the decommission process failed.

Tip: If you want to determine the status of a pending decommission process, issue the following command:

```
query process
```

3. Review the command output:
 - If status is provided for the decommission process, the process is in progress. For example:

```
query process
```

Process Number	Process Description	Process Status
-----	-----	-----

- If no status is provided for the decommission process, and you did not receive an error message, the process is incomplete. A process can be incomplete if files that are associated with the node are not yet deactivated. After the files are deactivated, run the decommission process again.
- If no status is provided for the decommission process, and you receive an error message, the process failed. Run the decommission process again.

Related reference:

[DECOMMISSION NODE \(Decommission a client node\)](#)

[DECOMMISSION VM \(Decommission a virtual machine\)](#)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Spectrum Protect™ server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

About this task

Some application clients always save data to the server as active backup data. Because active backup data is not managed by inventory expiration policies, the data is not deleted automatically, and uses server storage space indefinitely. To free the storage space that is used by obsolete data, you can deactivate the data.

When you run the deactivation process, all active backup data that was stored before the specified date becomes inactive. The data is deleted as it expires and cannot be restored. The deactivation feature applies only to application clients that protect Oracle databases.

Procedure

1. From the Operations Center Overview page, click Clients.
2. In the Clients table, select one or more clients and click More > Clean Up.
Command-line method: Deactivate data by using the DEACTIVATE DATA command.

Related reference:

[DEACTIVATE DATA \(Deactivate data for a client node\)](#)

Managing data storage

Manage your data for efficiency and add supported devices and media to the server to store client data.

- **Auditing a storage pool container**
Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.
- **Managing inventory capacity**
Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.
- **Managing memory and processor usage**
Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.
- **Tuning scheduled activities**
Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Related reference:

[Storage pool types](#)

Auditing a storage pool container

Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.

About this task

You audit a storage pool container in the following situations:

- When you issue the QUERY DAMAGED command and a problem is detected
- When the server displays messages about damaged data extents
- Your hardware reports an issue and error messages that are associated with the storage pool container are displayed

Procedure

1. To audit a storage pool container, issue the AUDIT CONTAINER command. For example, issue the following command to audit a container, 000000000000076c.dcf:

```
audit container c:\tsm-storage\07\000000000000076c.dcf
```

2. Review the output from the ANR4891I message for information about any damaged data extents.

What to do next

If you detect problems with the storage pool container, you can restore data based on your configuration. Issue the AUDIT CONTAINER command and specify the container name.

Related reference:

- [AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)
- [QUERY DAMAGED](#) (Query damaged data in a directory-container or cloud-container storage pool)

Managing inventory capacity

Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.

Before you begin

The active and archive logs have the following characteristics:

- The active log can be a maximum size of 512 GB. For more information about sizing the active log for your system, see [Planning the storage arrays](#).
- The archive log size is limited to the size of the file system that it is installed on. The archive log size is not maintained at a predefined size like the active log. Archive log files are automatically deleted after they are no longer needed.

As a best practice, you can optionally create an archive failover log to store archive log files when the archive log directory is full.

Check the Operations Center to determine the component of the inventory that is full. Ensure that you stop the server before you increase the size of one of the inventory components.

Procedure

- To increase the size of the database, complete the following steps:
 - Create one or more directories for the database on separate drives or file systems.
 - Issue the EXTEND DBSPACE command to add the directory or directories to the database. The directories must be accessible to the instance user ID of the database manager. By default, data is redistributed across all database directories and space is reclaimed.
 - Tips:
 - The time that is needed to complete redistribution of data and reclaiming of space is variable, depending on the size of your database. Make sure that you plan adequately.
 - Ensure that the directories that you specify are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.
 - Halt and restart the server to fully use the new directories.
 - Reorganize the database if necessary. Index and table reorganization for the server database can help to avoid unexpected database growth and performance issues. For more information about reorganizing the database, see [technote 1683633](#).

- To decrease the size of the database for V7.1 servers and later, issue the following DB2® commands from the server instance directory:
Restriction: The commands can increase I/O activity, and might affect server performance. To minimize performance problems, wait until one command is completed before you issue the next command. The DB2 commands can be issued when the server is running.

```
db2 connect to tsmdb1
db2 set schema tsmdb1
db2 ALTER TABLESPACE USERSPACE1 REDUCE MAX
db2 ALTER TABLESPACE IDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGEIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGESPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE5 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE5 REDUCE MAX
```

- To increase or decrease the size of the active log, complete the following steps:
 - Ensure that the location for the active log has enough space for the increased log size. If a log mirror exists, its location must also have enough space for the increased log size.
 - Halt the server.
 - In the dsmserv.opt file, update the ACTIVELOGSIZE option to the new size of the active log, in megabytes. The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

To change the active log to its maximum size of 512 GB, enter the following server option:

```
activelogsize 524288
```

- If you plan to use a new active log directory, update the directory name that is specified in the ACTIVELOGDIRECTORY server option. The new directory must be empty and must be accessible to the user ID of the database manager.
 - Restart the server.
- Compress the archive logs to reduce the amount of space that is required for storage. Enable dynamic compression of the archive log by issuing the following command:

```
setopt archlogcompress yes
```

Restriction: Use caution when you enable the ARCHLOGCOMPRESS server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log directory file system usage nears out of space conditions, the ARCHLOGCOMPRESS server option must be disabled. You can use the SETOPT command to disable archive log compression immediately without halting the server.

Related reference:

- 🔗 [ACTIVELOGSIZE server option](#)
- 🔗 [EXTEND DBSPACE \(Increase space for the database\)](#)
- 🔗 [SETOPT \(Set a server option for dynamic update\)](#)

Managing memory and processor usage

Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.

Before you begin

- Ensure that your configuration uses the required hardware and software. For more information, see [IBM Spectrum Protect™ Supported Operating Systems](#).
- For more information about managing resources such as the database and recovery log, see [Planning the storage arrays](#).
- Add more system memory to determine whether there is a performance improvement. Monitor memory usage regularly to determine whether more memory is required.

Procedure

1. Release memory from the file system cache where possible.
2. To manage the system memory that is used by each server on a system, use the `DBMEMPERCENT` server option. Limit the percentage of system memory that can be used by the database manager of each server. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.
3. Set the user data limit and private memory for the database to ensure that private memory is not exhausted. Exhausting private memory can result in errors, less than optimal performance, and instability.

Tuning scheduled activities

Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Procedure

1. Monitor system performance regularly to ensure that backup and maintenance tasks are completing successfully. For more information about monitoring, see [Monitoring a single-site disk solution](#).
2. If the monitoring information shows that the server workload increased, you might need to review the planning information. Review whether the capacity of the system is adequate in the following cases:
 - The number of clients increases
 - The amount of data that is being backed up increases
 - The amount of time that is available for backups changes
3. Determine whether your solution has performance issues. Review the client schedules to check whether tasks are completing within the scheduled time frame:
 - a. On the Clients page of the Operations Center, select the client.
 - b. Click Details.
 - c. From the client Summary page, review the Backed up and Replicated activity to identify any risks. Adjust the time and frequency of client backup operations, if necessary.
4. Schedule enough time for the following maintenance tasks to complete successfully within a 24-hour period:
 - a. Back up the database
 - b. Run expiration to remove client backups and archive file copies from server storage.

Related concepts:

- 🔗 [Performance](#)

Related tasks:

- 🔗 [Deduplicating data \(V7.1.1\)](#)

Securing the IBM Spectrum Protect server

Secure the IBM Spectrum Protect™ server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.

- Security concepts
You can protect IBM Spectrum Protect from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.
- Managing administrators
An administrator who has system authority can complete any task with the IBM Spectrum Protect server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.
- Changing password requirements
You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect.
- Securing the server on the system
Protect the system where the IBM Spectrum Protect server runs to prevent unauthorized access.

Security concepts

You can protect IBM Spectrum Protect™ from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.

Transport Layer Security

You can use the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol to provide transport layer security for a secure connection between servers, clients, and storage agents. If you send data between the server, client, and storage agent, use SSL or TLS to encrypt the data.

Tip: Any IBM Spectrum Protect documentation that indicates "SSL" or to "select SSL" applies to TLS.

SSL is provided by the Global Security Kit (GSKit) that is installed with the IBM Spectrum Protect server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with a DB2® database instance that is used by any IBM Spectrum Protect servers.

Each server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a certificate authority (CA). You can use your own certificates or purchase certificates from a CA. Either certificate must be installed and added to the key database on the IBM Spectrum Protect server, client, or storage agent. The certificate is verified by the SSL client or server that requests or initiates the SSL communication. Some CA certificates are preinstalled in the key databases, by default.

SSL is set up independently on the IBM Spectrum Protect server, client, and storage agent.

Authority levels

With each IBM Spectrum Protect server, different administrative authority levels are available that determine the tasks that an administrator can complete.

After registration, an administrator must be granted authority by being assigned one or more administrative authority levels. An administrator with system authority can complete any task with the server and assign authority levels to other administrators by using the GRANT AUTHORITY command. Administrators with policy, storage, or operator authority can complete subsets of tasks.

An administrator can register other administrator IDs, grant levels of authority to them, rename IDs, remove IDs, and lock and unlock them from the server.

An administrator can control access to specific client nodes for root user IDs and non-root user IDs. By default, a non-root user ID cannot back up data on the node. Use the UPDATE NODE command to change the node settings to enable backup.

Passwords

By default, the server automatically uses password authentication. With password authentication, all users must enter a password when they access the server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see [Managing passwords and logon procedures \(V7.1.1\)](#).

Table 1. Password authentication characteristics

Characteristic	More information
Case-sensitivity	Not case-sensitive.
Default password expiration	90 days. The expiration period begins when an administrator ID or client node is first registered to the server. If the password is not changed within this period, the password must be changed the next time that the user accesses the server.
Invalid password attempts	You can set a limit on consecutive invalid password attempts for all client nodes. When the limit is exceeded, the server locks the node.
Default password length	8 characters. The administrator can specify a minimum length. Beginning with Version 8.1.4, the default minimum length for server passwords changed from 0 to 8 characters.

Session security

Session security is the level of security that is used for communication among IBM Spectrum Protect client nodes, administrative clients, and servers and is set by using the SESSIONSECURITY parameter.

The SESSIONSECURITY parameter can be set to one of the following values:

- The STRICT value enforces the highest level of security for communication between IBM Spectrum Protect servers, nodes, and administrators.
- The TRANSITIONAL value specifies that the existing communication protocol is used while you update your IBM Spectrum Protect software to V8.1.2 or later. This is the default. When SESSIONSECURITY=TRANSITIONAL, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

Note: You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Spectrum Protect servers, but you are not required to upgrade servers first. Communication between servers and clients is not interrupted.

For more information about the SESSIONSECURITY parameter values, see the following commands.

Table 2. Commands used to set the SESSIONSECURITY parameter

Entity	Command
Client nodes	<ul style="list-style-type: none"> • REGISTER NODE • UPDATE NODE
Administrators	<ul style="list-style-type: none"> • REGISTER ADMIN • UPDATE ADMIN
Servers	<ul style="list-style-type: none"> • DEFINE SERVER • UPDATE SERVER

Administrators that authenticate by using the DSMADMC command, DSMC command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

Tips:

- Ensure that all IBM Spectrum Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.
- After an administrator authenticates to a V8.1.2 or later server by using a V8.1.2 or later client, the administrator can authenticate only on clients or servers that are using V8.1.2 or later. An administrator command can be issued from any system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Enforce the highest level of security for communication with the IBM Spectrum Protect server by ensuring that all nodes, administrators, and servers use STRICT session security. You can use the SELECT command to determine which servers, nodes, and administrators are using TRANSITIONAL session security and should be updated to use STRICT session security.

Related tasks:

[↗ Securing communications](#)

Managing administrators

An administrator who has system authority can complete any task with the IBM Spectrum Protect™ server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.

Procedure

Complete the following tasks to modify administrator settings.

Task	Procedure
Add an administrator.	<p>To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:</p> <ol style="list-style-type: none"> Register the administrator and specify Pa\$#\$twO as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$twO</pre> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>
Change administrative authority.	<p>Change the authority level for an administrator, ADMIN1.</p> <ul style="list-style-type: none"> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre> Revoke system authority for the administrator by issuing the following command: <pre>revoke authority admin1 classes=system</pre>
Remove administrators.	<p>Remove an administrator, ADMIN1, from accessing the IBM Spectrum Protect server by issuing the following command:</p> <pre>remove admin admin1</pre>
Temporarily prevent access to the server.	<p>Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.</p>

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect™.

About this task

By enforcing password authentication and managing password restrictions, you protect your data and your servers from potential security risks.

Procedure

Complete the following tasks to change password requirements for IBM Spectrum Protect servers.

Table 1. Authentication tasks for IBM Spectrum Protect servers

Task	Procedure
Set a limit for invalid password attempts.	<ol style="list-style-type: none">On the Servers page in the Operations Center, select the server.Click Details, and then click the Properties tab.Set the number of invalid attempts in the Invalid sign-on attempt limit field. <p>The default value at installation is 0.</p>
Set a minimum length for passwords.	<ol style="list-style-type: none">On the Servers page in the Operations Center, select the server.Click Details and then click the Properties tab.Set the number of characters in the Minimum password length field.
Set the expiration period for passwords.	<ol style="list-style-type: none">On the Servers page in the Operations Center, select the server.Click Details and then click the Properties tab.Set the number of days in the Password common expiration field.
Disable password authentication.	<p>By default, the server automatically uses password authentication. With password authentication, all users must enter a password to access the server.</p> <p>You can disable password authentication only for passwords that authenticate with the server (LOCAL). By disabling password authentication, you increase the security risk for the server.</p>
Set a default authentication method.	<p>Issue the SET DEFAULTAUTHENTICATION command. For example, to use the server as the default authentication method, issue the following command:</p> <pre>set defaultauthentication local</pre> <p>To update one client node to authenticate with the server, include AUTHENTICATION=LOCAL in the UPDATE NODE command:</p> <pre>update node authentication=local</pre>

Related concepts:

- [Authenticating IBM Spectrum Protect users by using an LDAP server](#)
- [Managing passwords and logon procedures \(V7.1.1\)](#)

Securing the server on the system

Protect the system where the IBM Spectrum Protect™ server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the directories for the server database and the server instance. Keep the access settings for these directories that you configured during implementation.

- Restricting user access to the server
Authority levels determine what an administrator can do with the IBM Spectrum Protect server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.
- Limiting access through port restrictions
Limit access to the server by applying port restrictions.

Restricting user access to the server

Authority levels determine what an administrator can do with the IBM Spectrum Protect™ server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.

Procedure

1. After you register an administrator by using the REGISTER ADMIN command, use the GRANT AUTHORITY command to set the administrator's authority level. For details about setting and changing authority, see Managing administrators.
2. To control the authority of an administrator to complete some tasks, use the following two server options:
 - a. You can select the authority level that an administrator must have to issue QUERY and SELECT commands with the QUERYAUTH server option. By default, no authority level is required. You can change the requirement to one of the authority levels, including system.
 - b. You can specify that system authority is required for commands that cause the server to write to an external file with the REQSYSAUTHOUTFILE server option. By default, system authority is required for such commands.
3. You can restrict data backup on a client node to only root user IDs or authorized users. For example, to limit backups to the root user ID, issue the REGISTER NODE or UPDATE NODE command and specify the BACKUPINITIATION=root parameter:

```
update node backupinitiation=root
```

Limiting access through port restrictions

Limit access to the server by applying port restrictions.

About this task

You might have to restrict access to specific servers, based on your security requirements. The IBM Spectrum Protect™ server can be configured to listen on four TCP/IP ports: two that can be used for either regular TCP/IP protocols or Secure Sockets Layer (SSL)/Transport Layer Security (TLS) protocols and two that can be used only for the SSL/TLS protocol.

Procedure

You can set the server options to specify the port that you require, as listed in Table 1.

Table 1. Server options and port access

Server option	Port access
TCPPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default value is 1500.
TCPADMINPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default is the value of TCPPORT. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.
SSLTCPPORT	Specifies the SSL TCP/IP port address for a server. This port listens for SSL-enabled sessions only. A default port value is not available.

Server option	Port access
SSLTCPADMINPORT	<p>Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. A default port value is not available.</p> <p>Use this option to separate administrative client traffic from regular client traffic that uses the TCPSPORT and SSLTCPSPORT options.</p>

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (SSLTCPSPORT and SSLTCPADMINPORT):

- When you specify the server's SSL-only port for the LLADDRESS on the DEFINE SERVER or UPDATE SERVER command, you must also specify the SSL=YES parameter.
- When you specify the server's SSL-only port for the client's TCPSPORT option, you must also specify YES for the SSL client option.

Related reference:

Planning firewall access

Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.

Before you begin

You must have system or operator privilege to stop and start the IBM Spectrum Protect™ server.

- **Stopping the server**
Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.
- **Starting the server for maintenance or reconfiguration tasks**
Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Stopping the server

Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.

About this task

When you issue the HALT command to stop the server, the following actions occur:

- All processes and client node sessions are canceled.
- All current transactions are stopped. (The transactions will be rolled back when the server is restarted.)

Procedure

To prepare the system and stop the server, complete the following steps:

1. Prevent new client node sessions from starting by issuing the DISABLE SESSIONS command:

```
disable sessions all
```

2. Determine whether any client node sessions or processes are in progress by completing the following steps:
 - a. On the Overview page of the Operations Center, view the Activity area for the total numbers of processes and sessions that are currently active. If numbers differ significantly from the usual numbers that are displayed during your daily storage-management routine, view other status indicators in the Operations Center to check whether there is a problem.
 - b. View the graph in the Activity area to compare the amount of network traffic over the following periods:
 - The current period, that is, the most recent 24-hour period

- The previous period, that is, the 24 hours before the current period
- If the graph for the previous period represents the expected amount of traffic, significant differences on the graph for the current period might indicate a problem.
- c. On the Servers page, select a server for which you want to view processes and sessions, and click Details. If the server is not registered as a hub or spoke server in the Operations Center, obtain information about processes by using administrative commands. Issue the QUERY PROCESS command to query processes and obtain information about sessions by issuing the QUERY SESSION command.
3. Wait until the client node sessions are completed or cancel them. To cancel processes and sessions, complete the following steps:
 - o On the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - o Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - o Click Cancel.
 - o If the server is not registered as a hub or spoke server in the Operations Center, cancel sessions by using administrative commands. Issue the CANCEL SESSION command to cancel a session and cancel processes by using the CANCEL PROCESS command.

Tip: If the process that you want to cancel is waiting for a tape volume to be mounted, the mount request is canceled. For example, if you issue an EXPORT, IMPORT, or MOVE DATA command, the command might initiate a process that requires a tape volume to be mounted. However, if a tape volume is being mounted by an automated library, the cancel operation might not take effect until the mount process is complete. Depending on your system environment, this could take several minutes.
 4. Stop the server by issuing the HALT command:

```
halt
```

Starting the server for maintenance or reconfiguration tasks

Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSEV utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, dsmserv.opt, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see Starting a server in maintenance mode.

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

halt

2. Start the server by using the method that you use in production mode. Follow the instructions for your operating system:
 - o **AIX** Starting the server instance
 - o **Linux** Starting the server instance
 - o **Windows** Starting the server instance

Operations that were disabled during maintenance mode are reenabled.

Planning to upgrade the server

When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect™ server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.

About this task

Follow these guidelines:

- The preferred method is to upgrade the server by using the installation wizard. After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.
- If upgrades are available for both the server component and the Operations Center component, select the check boxes to upgrade both components.

Procedure

1. Review the list of fix packs and interim fixes. See technote 1239415.
2. Review product improvements, which are described in readme files.
Tip: When you obtain the installation package file from the IBM Spectrum Protect support site, you can also access the readme file.
3. Ensure that the version that you upgrade your server to is compatible with other components, such as storage agents and library clients. See technote 1302789.
4. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted. See technote 1053218.
5. Review the upgrade instructions. Ensure that you back up the server database, the device configuration information, and the volume history file.

What to do next

To install a fix pack or interim fix, follow the instructions for your operating system:

- **AIX** Installing an IBM Spectrum Protect server fix pack
- **Linux** Installing an IBM Spectrum Protect server fix pack
- **Windows** Installing an IBM Spectrum Protect server fix pack

Related information:

[Upgrade and Migration Process - Frequently Asked Questions](#)

Preparing for an outage or system update

Prepare IBM Spectrum Protect™ to maintain your system in a consistent state during a planned power outage or system update.

About this task

Ensure that you schedule activities regularly to manage, protect, and maintain the server.

Procedure

1. Cancel processes and sessions that are in progress by completing the following steps:
 - a. In the Operations Center, on the Servers page, select a server for which you want to view processes and sessions, and click Details.

- b. Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - c. Click Cancel.
2. Stop the server by issuing the HALT command:

```
halt
```

Tip: You can issue the halt command from the Operations Center by hovering over the Settings icon and clicking Command Builder. Then, select the server, type `halt`, and press Enter.

Implementing a disaster recovery plan

Implement a disaster recovery strategy to recover your applications if a disaster occurs and to ensure high server availability.

About this task

Determine your disaster recovery requirements by identifying the business priorities for client node recovery, the systems that you use to recover data, and whether client nodes have connectivity to a recovery server. Use replication and storage pool protection to protect data. You must also determine how often directory-container storage pools are protected.

- **Completing recovery drills**
Schedule disaster recovery drills to prepare for audits that certify the recoverability of the IBM Spectrum Protect server and to ensure that data can be restored and operations can resume after an outage. A drill also helps you ensure that all data can be restored and operations resumed before a critical situation occurs.

Recovering from system outages

For IBM Spectrum Protect™ single-site disk solutions, you can recover the inventory locally only and restore the database to protect your data.

Procedure

Use one of the following methods to recover inventory to a local site, based on the type of information that is backed up. Restriction: Because single-site disk solutions do not have a second copy of the storage pool, you cannot restore storage pools. To review the architecture of disk solutions, see [Selecting an IBM Spectrum Protect solution for your environment](#).

Table 1. Scenarios for recovering from a disaster

Scenario	Procedure
Your system is inaccessible and you want to locally restore to an earlier version by using system tools.	<ul style="list-style-type: none"> • Use IBM Spectrum Protect to back up the server to another server. • Use operating system tools to back up and restore your system to an earlier version.
An outage or disaster occurs and you want to restore your data from backed up versions of the data.	<ul style="list-style-type: none"> • To back up a client, on the TSM Clients page of the Operations Center, select the clients that you want to back up, and click Back Up. • On the TSM Servers page of the Operations Center, select the server whose database you want to back up. Click Back Up, and follow the instructions in the Back Up Server Database window. <p>To restore a storage pool from a backed-up version of the storage pool, you must restore the database. Issue the <code>DSMSERV RESTORE DB</code> command to restore the database and associated storage pools to a backed-up version.</p>

- **Restoring the database**
You might have to restore the IBM Spectrum Protect database after a disaster. You can restore the database to the most

current state or to a specified point in time. You must have full, incremental, or snapshot database backup volumes to restore the database.

Related reference:

- [AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)
- [DSMSERV RESTORE DB](#) (Restore the database)

Multisite disk solution

This data protection solution provides replication at multiple sites so that each server protects data for the other site.

- **Planning for a multisite disk data protection solution**
Plan for a multisite disk data protection solution with servers at two sites that use data deduplication and replication.
- **Multisite disk implementation of a data protection solution**
The multisite disk solution is configured at two sites and uses data deduplication and replication.
- **Monitoring a multisite disk solution**
After you implement a multisite disk solution with IBM Spectrum Protect, monitor the solution to ensure correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.
- **Managing operations for a multisite disk solution**
Use this information to manage operations for a multisite disk solution with IBM Spectrum Protect that includes a server and uses data deduplication for multiple locations.

Planning for a multisite disk data protection solution

Plan for a multisite disk data protection solution with servers at two sites that use data deduplication and replication.

Implementation methods

You can configure servers for a multisite disk solution in the following ways:

Configure servers by using the Operations Center and administrative commands

You can configure a range of storage systems and the server software for your solution. Configuration tasks are completed by using wizards and options in the Operations Center and IBM Spectrum Protect™ commands. For information about getting started, see the Planning roadmap.

Configure the servers by using automated scripts


For detailed guidance on configuration with specific IBM® Storwize® storage systems, and by using automated scripts to configure each server, see the IBM Spectrum Protect blueprints. The documentation and scripts are available on IBM developerWorks® at IBM Spectrum Protect Blueprints.

The blueprint documentation does not include steps for installing and configuring the Operations Center, or setting up secure communications by using Transport Security Layer (TLS). Replication is configured by using commands after each server is set up. An option for using Elastic Storage Server, based on IBM Spectrum Scale™ technology, is included.

Planning roadmap

Plan for a multisite disk solution by reviewing the architecture layout in the following figure and then completing the roadmap tasks that follow the diagram.

Figure 1. Multisite disk solution



The following steps are required to plan properly for a multisite disk environment.

1. Select your system size.
2. Plan for the sites.
3. Meet system requirements for hardware and software.
4. Record values for your system configuration in the planning worksheets.
5. Plan for storage.
6. Plan for security.
 - a. Plan for administrator roles.
 - b. Plan for secure communications.
 - c. Plan for storage of encrypted data.
 - d. Plan for firewall access.

Selecting a system size

Select the size of the IBM Spectrum Protect™ server based on the amount of data that you manage and the systems to be protected.

About this task

You can use the information in the table to determine the size of the server that is required, based on the amount of data that you manage.

The following table describes the volume of data that a server manages. This amount includes all versions. The daily amount of data is how much new data you back up each day. Both the total managed data and daily amount of new data are measured as the size before any data reduction.

Table 1. Determining the size of the server

Total managed data	Daily amount of new data to back up	Required server size
60 TB - 240 TB	Up to 10 TB per day	Small
196 TB - 784 TB	10 - 20 TB per day	Medium
1000 TB - 4000 TB	20 - 100 TB per day	Large

The daily backup values in the table are based on test results with 128 MB sized objects, which are used by IBM Spectrum Protect for Virtual Environments. Workloads that consist of objects that are smaller than 128 KB might not be able to achieve these daily limits.

Planning the sites

Review use cases and evaluate the factors to provide the most efficient data protection for the multisite disk solution for IBM Spectrum Protect™.

Use cases

The multisite disk solution creates at least one copy of backed-up data. If the IBM Spectrum Protect servers are at separate locations, the backed-up replica is maintained offsite.

Tip: Avoid conflicts in managing administrative IDs and client option sets by identifying the IDs and option sets that will be replicated to the target server and the IDs and option sets that will be managed in an enterprise configuration. You cannot define an administrative user ID for a registered node if an administrative ID exists for the same node.

Although your company might benefit from a multisite disk solution for various reasons, the most common reasons to use a multisite disk solution include the following replication scenarios:

Replication from the primary site to the disaster recovery site

In this scenario, data that is backed up from the primary site, Site A, is replicated to a server at the secondary, disaster recovery site, Site B. If a disaster occurs at Site A, such as failure of the server, you can use the server at Site B to recover systems. Alternatively, you can use the server at Site A to restore primary storage pool data at Site B, such as after a disk storage failure at Site B.

Mutual replication at two active sites

In this scenario, local data at each site is backed up by the servers at both Site A and Site B. Data that is backed up from Site A is replicated to Site B, and backed-up data from Site B is replicated to Site A. If data that was backed up is lost at Site A, you can use the server at Site B to recover storage pool data to the server at Site A. If Site A is no longer available, you can recover the replicated data for Site A to a new system at Site B. You must size the server resources to ensure that either server has sufficient capacity to back up and restore all client nodes as part of your disaster recovery plan.

Protect remote servers to the primary site

In this scenario, you configure remote servers that are relatively small to replicate data that is backed up to a larger server at the primary site. If bandwidth is limited, it might not be practical to restore systems to the remote sites. In this case, you might want to recover systems at the primary site before you replicate the backed-up data to the remote servers.

Factors to evaluate

Before you implement a multisite disk solution, evaluate the following factors:

Network bandwidth

The network must have sufficient bandwidth for the expected data transfers between nodes, for replication, and for the cross-site restore operations that are required for disaster recovery. Before you proceed with testing replication throughput, ensure that your network can handle the replication traffic. Calculate the required network bandwidth for the steady-state requirement by applying the guidelines in Estimating network bandwidth required for replication (V7.1.1).

The network connection is often a shared resource. Plan the time of day to schedule node replication to run to avoid a conflict with other resource users. Also, network controls might limit activity to only a portion of the bandwidth. There are no controls in IBM Spectrum Protect to restrict network usage.

Resources for the initial replication

To set up the data protection solution across two sites, you must replicate data initially from Site A to the target server at Site B. To ensure that the initial replication is successful, you must determine whether you have the network bandwidth, processor resources, and time available to replicate the data. You might have to plan for replicating the initial full backups across several days. If you cannot extend the schedule for the initial backups, you can replicate data from Site A to Site B without using the network. For example, you can export and import the backed-up data by using media or you can temporarily locate the source and target servers on the same site.

Daily data ingestion

For the multisite disk solution, the daily data ingestion and total data retention must be within the capacity of the configurations. For example, a large configuration has a data ingestion capacity of up to 100 TB per day, including node replication. In cases where the backup requirements exceed the capacity of a single server, you can configure a solution that uses multiple servers to achieve the required capacity.

Server configuration

The server configuration must meet or exceed the requirements for the multisite disk solution.

Single replica of backed-up data

The multisite disk solution is most efficient when a single, offsite copy of the backed-up data meets your data protection and risk mitigation requirements. In this case, the single copy of the data is maintained off-site at the location of a replication server.

Related reference:

System requirements for a multisite disk solution

System requirements for a multisite disk solution

After you select the IBM Spectrum Protect™ solution that best fits your data protection requirements, review the system requirements to plan for implementation of the data protection solution.

Ensure that your system meets the hardware and software prerequisites for the size of server that you plan to use.

- **Hardware requirements**
Hardware requirements for your IBM Spectrum Protect solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.
- **Software requirements**
Documentation for the IBM Spectrum Protect multisite disk solution includes installation and configuration tasks for the following operating systems. You must meet the minimum software requirements that are listed.

Related information:

[IBM Spectrum Protect Supported Operating Systems](#)

Hardware requirements

Hardware requirements for your IBM Spectrum Protect™ solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.

For a definition of system sizes, see [t_msdisk_select_size.html](#).

The following table includes minimum hardware requirements for the server and storage, based on the size of the server that you plan to build. If you are using local partitions (LPARs) or work partitions (WPARs), adjust the network requirements to take account of the partition sizes.

Use the information in the following table as a starting point. For the most up-to-date information about hardware requirements and specifications for the server and storage, see the IBM Spectrum Protect Blueprints.

Hardware component	Small system	Medium system	Large system												
Server processor	<table border="0"><tr><td>AIX</td><td>6 processor cores, 3.42 GHz or faster</td></tr><tr><td>Linux Windows</td><td>16 processor cores, 1.7 GHz or faster</td></tr></table>	AIX	6 processor cores, 3.42 GHz or faster	Linux Windows	16 processor cores, 1.7 GHz or faster	<table border="0"><tr><td>AIX</td><td>10 processor cores, 3.42 GHz or faster</td></tr><tr><td>Linux Windows</td><td>20 processor cores, 2.2 GHz or faster</td></tr></table>	AIX	10 processor cores, 3.42 GHz or faster	Linux Windows	20 processor cores, 2.2 GHz or faster	<table border="0"><tr><td>AIX</td><td>20 processor cores, 3.42 GHz</td></tr><tr><td>Linux Windows</td><td>44 processor cores, 2.2 GHz or faster</td></tr></table>	AIX	20 processor cores, 3.42 GHz	Linux Windows	44 processor cores, 2.2 GHz or faster
AIX	6 processor cores, 3.42 GHz or faster														
Linux Windows	16 processor cores, 1.7 GHz or faster														
AIX	10 processor cores, 3.42 GHz or faster														
Linux Windows	20 processor cores, 2.2 GHz or faster														
AIX	20 processor cores, 3.42 GHz														
Linux Windows	44 processor cores, 2.2 GHz or faster														
Server memory	64 GB RAM	128 GB RAM	256 GB RAM												

Hardware component	Small system	Medium system	Large system
Network	<ul style="list-style-type: none"> • 10 GB Ethernet (1 port) • 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> • 10 GB Ethernet (2 ports) • 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> • 10 GB Ethernet (4 ports) • 8 GB Fibre Channel adapter (4 ports)
Storage	<ul style="list-style-type: none"> • 1.45 TB SSD disks for the database, plus space for Operations Center records • 67 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> • 2.53 TB SSD disks for the database, plus space for Operations Center records • 207.9 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> • 6.54 TB SSD disks for the database, plus space for Operations Center records • 1049.67 TB deduplicated directory-container storage pool

Implementing the correct processor core technology

You must use the correct type of processor core technology for the server processor. For information about the type of processor core technology, see the IBM Spectrum Protect Blueprints.

Estimating database space requirements for the Operations Center

Hardware requirements for the Operations Center are included in the preceding table, except for the database and archive log space (inventory) that the Operations Center uses to hold records for managed clients.

If you do not plan to install the Operations Center on the same system as the server, you can estimate system requirements separately. To calculate system requirements for the Operations Center, see the system requirements calculator in technote 1641684.

Managing the Operations Center on the server is a workload that requires extra space for database operations. The amount of space depends on the number of clients that are monitored on a server. Review the following guidelines to estimate how much space your server requires.

Database space

The Operations Center uses approximately 1.2 GB of database space for every 1000 clients that are monitored on a server. For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1500 clients. This configuration has a total of 6500 clients across the four servers and requires approximately 8.4 GB of database space. This value is calculated by rounding the 6500 clients up to the next closest 1000, which is 7000:

$$7 \times 1.2 \text{ GB} = 8.4 \text{ GB}$$

Archive log space

The Operations Center uses approximately 8 GB of archive log space every 24 hours, for every 1000 clients. In the example of 6500 clients across the hub server and the spoke servers, 56 GB of archive log space is used over a 24-hour period for the hub server.

For each spoke server in the example, the archive log space that is used over 24 hours is approximately 16 GB. These estimates are based on the default status collection interval of 5 minutes. If you reduce the collection interval from once every 5 minutes to once every 3 minutes, the space requirements increase. The following examples show the approximate increase in the log space requirement with a collection interval of once every 3 minutes:

- Hub server: 56 GB to approximately 94 GB
- Each spoke server: 16 GB to approximately 28 GB

Increase the archive log space so that you have sufficient space available to support the Operations Center, without affecting the existing server operations.

Hardware requirements for the second server

If you are planning to set up your sites so that everything at the first site is replicated to the second site, hardware requirements are identical at both sites. If you want to only replicate a subset of data to your second site, storage and network requirements might be reduced.

Software requirements

Documentation for the IBM Spectrum Protect™ multisite disk solution includes installation and configuration tasks for the following operating systems. You must meet the minimum software requirements that are listed.

For information about software requirements for IBM® lin_tape device drivers, refer to the IBM Tape Device Drivers Installation and User's Guide.

AIX systems

Type of software	Minimum software requirements
Operating system	IBM AIX® 7.1 For more information about operating system requirements, see AIX: Minimum system requirements for AIX systems.
Gunzip utility	The gunzip utility must be available on your system before you install or upgrade the IBM Spectrum Protect server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.
File system type	JFS2 file systems AIX systems can cache a large amount of file system data, which can reduce memory that is required for server and IBM DB2® processes. To avoid paging with the AIX server, use the rbrw mount option for the JFS2 file system. Less memory is used for the file system cache and more is available for IBM Spectrum Protect. Do not use the file system mount options, Concurrent I/O (CIO), and Direct I/O (DIO), for file systems that contain the IBM Spectrum Protect database, logs, or storage pool volumes. These options can cause performance degradation of many server operations. IBM Spectrum Protect and DB2 can still use DIO where it is beneficial to do so, but IBM Spectrum Protect does not require the mount options to selectively take advantage of these techniques.
Other software	Korn Shell (ksh)

Linux systems

Type of software	Minimum software requirements
Operating system	Red Hat Enterprise Linux 7 (x86_64)
Libraries	GNU C libraries, Version 2.3.3-98.38 or later that is installed on the IBM Spectrum Protect system. Red Hat Enterprise Linux Servers: <ul style="list-style-type: none">• libaio• libstdc++.so.6 (32-bit and 64-bit packages are required)• numactl.x86_64
File system type	Format database-related file systems with ext3 or ext4. For storage pool-related file systems, use XFS.
Other software	Korn Shell (ksh)

Windows systems

Type of software	Minimum software requirements
Operating system	Microsoft Windows Server 2012 R2 (64-bit) or Windows Server 2016
File system type	NTFS

Type of software	Minimum software requirements
Other software	<p>Windows 2012 R2 or Windows 2016 with .NET Framework 3.5 is installed and enabled.</p> <p>The following User Account Control policies must be disabled:</p> <ul style="list-style-type: none"> • User Account Control: Admin Approval Mode for the Built-in Administrator account • User Account Control: Run all administrators in Admin Approval Mode

Related tasks:

[Setting AIX network options](#)

Planning worksheets

Use the planning worksheets to record values that you use to set up your system and configure the IBM Spectrum Protect™ server. Use the best practice default values that are listed in the worksheets.

Each worksheet helps you prepare for different parts of the system configuration by using best practice values:

Server system preconfiguration

Use the preconfiguration worksheets to plan for the file systems and directories that you create when you configure file systems for IBM Spectrum Protect during system setup. All directories that you create for the server must be empty.

Server configuration

Use the configuration worksheets when you configure the server. Default values are suggested for most items, except where noted.



Table 1. Worksheet for preconfiguration of an AIX server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	<p>Ensure that this port is available when you install and configure the operating system</p> <p>The port number can be a number in the range 1024 - 32767.</p>
Directory for the server instance	/home/tsminst1/tsminst1		50 GB	If you change the value for the server instance directory from the default, also modify the DB2® instance owner value in Table 2.
Directory for server installation	/		Available space that is required for the directory: 5 GB	
Directory for server installation	/usr		Available space that is required for the directory: 5 GB	
Directory for server installation	/var		Available space that is required for the directory: 5 GB	

Item	Default value	Your value	Minimum directory size	Notes
Directory for server installation	/tmp		Available space that is required for the directory: 5 GB	
Directory for server installation	/opt		Available space that is required for the directory: 10 GB	
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	When you create the active log during the initial configuration of the server, set the size to 128 GB.
Directory for the archive log	/tsminst1/TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 ...		<p>Minimum total space for all directories:</p> <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	<p>Create a minimum number of file systems for the database, depending on the size of your system:</p> <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		<p>Minimum total space for all directories:</p> <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	<p>Create a minimum number of file systems for storage, depending on the size of your system:</p> <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	<p>Create a minimum number of file systems for backing up the database, depending on the size of your system:</p> <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 <p>The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.</p>

Table 2. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 1 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	passw0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Item	Default value	Your value	Notes
Schedule start time	22:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Linux

Table 3. Worksheet for preconfiguration of a Linux server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	<p>Ensure that this port is available when you install and configure the operating system</p> <p>The port number can be a number in the range 1024 - 32767.</p>
Directory for the server instance	/home/tsminst1/tsminst1		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 4.
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	
Directory for the archive log	/tsminst1/TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	

Item	Default value	Your value	Minimum directory size	Notes
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	Create a minimum number of file systems for the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 <p>The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.</p>

Table 4. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
------	---------------	------------	-------

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 3 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	passw0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.
Schedule start time	22:00		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window. Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Windows

Because many volumes are created for the server, configure the server by using the Windows feature of mapping disk volumes to directories rather than to drive letters.

For example, C:\tsminst1\TSMdbpsace00 is a mount point to a volume with its own space. The volume is mapped to a directory under the C: drive, but does not take up space from the C: drive. The exception is the server instance directory, C:\tsminst1, which can be a mount point or a regular directory.

Table 5. Worksheet for preconfiguration of a Windows server system

Item	Default value	Your value	Minimum directory size	Notes
------	---------------	------------	------------------------	-------

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	C:\tsminst1		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 6.
Directory for the active log	C:\tsminst1\TSMalog		<ul style="list-style-type: none"> • Small and medium: 140 GB • Large: 300 GB 	
Directory for the archive log	C:\tsminst1\TSMarchlog		<ul style="list-style-type: none"> • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	C:\tsminst1\TSMdbspace00 C:\tsminst1\TSMdbspace01 C:\tsminst1\TSMdbspace02 C:\tsminst1\TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	Create a minimum number of file systems for the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	C:\tsminst1\TSMbkup00 C:\tsminst1\TSMbkup01 C:\tsminst1\TSMbkup02 C:\tsminst1\TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems, but preferably 6 The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Table 6. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 5 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	pAssW0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Item	Default value	Your value	Notes
Schedule start time	22:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Planning for storage

Choose the most effective storage technology for IBM Spectrum Protect™ components to ensure efficient server performance and operations.

Storage hardware devices have different capacity and performance characteristics, which determine how they can be used effectively with IBM Spectrum Protect. For general guidance on selecting the appropriate storage hardware and set up for your solution, review the following guidelines.

Database and active log

- Use a fast disk for the IBM Spectrum Protect database and active log, for example with the following characteristics:
 - High performance, 15k rpm disk with Fibre Channel or serial-attached SCSI (SAS) interface
 - Solid-state disk (SSD)
- Isolate the active log from the database unless you use SSD or flash hardware
- When you create arrays for the database, use RAID level 5

Storage pool

- You can use less expensive and slower disks for the storage pool
- The storage pool can share disks for the archive log and database backup storage
- Use RAID level 6 for storage pool arrays to add protection against double drive failures when you use large disk types
- Planning the storage arrays

Prepare for disk storage configuration by planning for RAID arrays and volumes, according to the size of your IBM Spectrum Protect system.

Related reference:

[Storage system requirements and reducing the risk of data corruption](#)

Planning for security

Plan to protect the security of systems in the IBM Spectrum Protect™ solution with access and authentication controls, and consider encrypting data and password transmission.

For guidelines about protecting your storage environment against ransomware attacks, and recovering your storage environment if an attack occurs, see Protecting the storage environment against ransomware.

- Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect solution.
- Planning for secure communications

Plan for protecting communications among the IBM Spectrum Protect solution components.

- Planning for storage of encrypted data
Determine whether your company requires stored data to be encrypted, and choose the option that best suits your needs.
- Planning firewall access
Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect solution to work.

Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect™ solution.

You can assign one of the following levels of authority to administrators:

System

Administrators with system authority have the highest level of authority. Administrators with this level of authority can complete any task. They can manage all policy domains and storage pools, and grant authority to other administrators.

Policy

Administrators who have policy authority can manage all of the tasks that are related to policy management. This privilege can be unrestricted, or can be restricted to specific policy domains.

Storage

Administrators who have storage authority can allocate and control storage resources for the server.

Operator

Administrators who have operator authority can control the immediate operation of the server and the availability of storage media such as tape libraries and drives.

The scenarios in Table 1 provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

Table 1. Scenarios for administrator roles

Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	<ul style="list-style-type: none"> • System authority: 1 administrator ID
An administrator for multiple servers also manages the overall system. Several other administrators manage their own storage pools.	<ul style="list-style-type: none"> • System authority on all servers: 1 administrator ID for the overall system administrator • Storage authority for designated storage pools: 1 administrator ID for each of the other administrators
An administrator manages 2 servers. Another person helps with the administration tasks. Two assistants are responsible for helping to ensure that important systems are backed up. Each assistant is responsible for monitoring the scheduled backups on one of the IBM Spectrum Protect servers.	<ul style="list-style-type: none"> • System authority on both servers: 2 administrator IDs • Operator authority: 2 administrator IDs for the assistants with access to the server that each person is responsible for

Planning for secure communications

Plan for protecting communications among the IBM Spectrum Protect™ solution components.

Determine the level of protection that is required for your data, based on regulations and business requirements under which your company operates.

If your business requires a high level of security for passwords and data transmission, plan on implementing secure communication with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocols.

TLS and SSL provide secure communications between the server and client, but can affect system performance. To improve system performance, use TLS for authentication without encrypting object data. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option for client-to-server communication, and the UPDATE SERVER=SSL parameter for server-to-server communication. Beginning in V8.1.2, TLS is used for authentication by default. If you decide to use TLS to encrypt entire sessions, use the protocol only for sessions where it is necessary and add processor resources on the server to manage the increase in network traffic. You can also try other options. For example, some networking devices such as routers and switches provide the TLS or SSL function.

You can use TLS and SSL to protect some or all of the different possible communication paths, for example:

- Operations Center: browser to hub; hub to spoke
- Client to server
- Server to server: node replication

Related tasks:

[🔗 Securing communications](#)

Planning for storage of encrypted data

Determine whether your company requires stored data to be encrypted, and choose the option that best suits your needs.

If your company requires the data in storage pools to be encrypted, then you have the option of using IBM Spectrum Protect™ encryption, or an external device such as tape for encryption.

If you choose IBM Spectrum Protect to encrypt the data, extra computing resources are required at the client that might affect the performance of backup and restore processes.

Related information:

[🔗 technote 1963635](#)

Planning firewall access

Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect™ solution to work.

Table 1 describes the ports that are used by the server, client, and Operations Center.

Table 1. Ports that are used by the server, client, and Operations Center

Item	Default	Direction	Description
Base port (TCPSPORT)	1500	Outbound/inbound	Each server instance requires a unique port. You can specify an alternative port number instead of using the default. The TCPSPORT option listens for both TCP/IP and SSL-enabled sessions from the client. For administrative client traffic, you can use the TCPADMINPORT and ADMINONCLIENTPORT options to set port values.
SSL-only port (SSLTCPSPORT)	No default	Outbound/inbound	This port is used if you want to restrict communication on the port to SSL-enabled sessions only. To support both SSL and non-SSL communications, use the TCPSPORT or TCPADMINPORT options.
SMB	45	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SSH	22	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SMTP	25	Outbound	This port is used to send email alerts from the server.
NDMP	No default	Inbound/outbound	<p>The server must be able to open an outbound NDMP control port connection to the NAS device. The outbound control port is the Low-Level Address in the data mover definition for the NAS device.</p> <p>During an NDMP filer-to-server restore, the server must be able to open an outbound NDMP data connection to the NAS device. The data connection port that is used during a restore can be configured on the NAS device.</p> <p>During NDMP filer-to-server backups, the NAS device must be able to open outbound data connections to the server and the server must be able to accept inbound NDMP data connections. You can use the server option NDMPPORTRANGE to restrict the set of ports available for use as NDMP data connections. You can configure a firewall for connections to these ports.</p>

Item	Default	Direction	Description
Replication	No default	Outbound/inbound	The port and protocol for the outbound port for replication are set by the DEFINE SERVER command that is used to set up replication. The inbound ports for replication are the TCP ports and SSL ports that the source server names in the DEFINE SERVER command.
Client schedule port	Client port: 1501	Outbound	The client listens on the port that is named and communicates the port number to the server. The server contacts the client if server prompted scheduling is used. You can specify an alternative port number in the client options file.
Long running sessions	KEEPALIVE setting: YES	Outbound	When the KEEPALIVE option is enabled, keepalive packets are sent during client-server sessions to prevent the firewall software from closing long-running, inactive connections.
Operations Center	HTTPS: 11090	Inbound	These ports are used for the Operations Center web browser. You can specify an alternative port number.
Client management service port	Client port: 9028	Inbound	The client management service port must be accessible from the Operations Center. Ensure that firewalls cannot prevent connections. The client management service uses the TCP port of the server for the client node for authentication by using an administrative session.

Multisite disk implementation of a data protection solution

The multisite disk solution is configured at two sites and uses data deduplication and replication.

Implementation roadmap

The following steps are required to set up a multisite disk environment.

1. Set up the system.
 - a. Configure the storage hardware and set up storage arrays for your environment size.
 - b. Install the server operating system.
 - c. Configure multipath I/O.
 - d. Create the user ID for the server instance.
 - e. Prepare file systems for IBM Spectrum Protect.
2. Install the server and Operations Center.
3. Configure the server and Operations Center.
 - a. Complete the initial configuration of the server.
 - b. Set server options.
 - c. Configure Secure Sockets Layer for the server and client.
 - d. Configure the Operations Center.
 - e. Register your IBM Spectrum Protect license.
 - f. Configure data deduplication.
 - g. Define data retention rules for your business.
 - h. Define server maintenance schedules.
 - i. Define client schedules.
4. Install and configure clients.
 - a. Register and assign clients to schedules.

Tip: Avoid conflicts in managing administrative IDs and client option sets by identifying the IDs and option sets that will be replicated to the target server and the IDs and option sets that will be managed in an enterprise configuration. You cannot define an administrative user ID for a registered node if an administrative ID exists for the same node.
 - b. Install and verify the client management service.
 - c. Configure the Operations Center to use the client management service.
5. Configure the second server.
 - a. Configure for SSL communication between the hub and spoke server.
 - b. Add the second server as a spoke.

- c. Enable replication.
6. Complete the implementation.

Setting up the system

To set up the system, you must first configure your disk storage hardware and the server system for IBM Spectrum Protect™.

- **Configuring the storage hardware**
To configure your storage hardware, review general guidance for disk systems and IBM Spectrum Protect.
- **Installing the server operating system**
Install the operating system on the server system and ensure that IBM Spectrum Protect server requirements are met. Adjust operating system settings as directed.
- **Configuring multipath I/O**
You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.
- **Creating the user ID for the server**
Create the user ID that owns the IBM Spectrum Protect server instance. You specify this user ID when you create the server instance during initial configuration of the server.
- **Preparing file systems for the server**
You must complete file system configuration for the disk storage to be used by the server.

Configuring the storage hardware

To configure your storage hardware, review general guidance for disk systems and IBM Spectrum Protect™.

Procedure

1. Provide a connection between the server and the storage devices by following these guidelines:
 - Use a switch or direct connection for Fibre Channel connections.
 - Consider the number of ports that are connected and account for the amount of bandwidth that is needed.
 - Consider the number of ports on the server and the number of host ports on the disk system that are connected.
2. Verify that device drivers and firmware for the server system, adapters, and operating system are current and at the recommended levels.
3. Configure storage arrays. Make sure that you planned properly to ensure optimal performance. See Planning for storage for more information.
4. Ensure that the server system has access to disk volumes that are created. Complete the following steps:
 - a. If the system is connected to a Fibre Channel switch, zone the server to see the disks.
 - b. Map all of the volumes to tell the disk system that this specific server is allowed to see each disk.

Related tasks:

[Configuring storage](#)

Installing the server operating system

Install the operating system on the server system and ensure that IBM Spectrum Protect™ server requirements are met. Adjust operating system settings as directed.

- **Installing on AIX systems**
Complete the following steps to install AIX® on the server system.
- **Installing on Linux systems**
Complete the following steps to install Linux x86_64 on the server system.
- **Installing on Windows systems**
Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect server.

Installing on AIX systems

Complete the following steps to install AIX® on the server system.

Procedure

1. Install AIX Version 7.1, TL4, SP2, or later according to the manufacturer instructions.
2. Configure your TCP/IP settings according to the operating system installation instructions.
3. Open the /etc/hosts file and complete the following actions:

- o Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7 server.yourdomain.com server
```

- o Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1 localhost
```

4. Enable AIX I/O completion ports by issuing the following command:

```
chdev -l iocp0 -P
```

Server performance can be affected by the Olson time zone definition.

5. To optimize performance, change your system time zone format from Olson to POSIX. Use the following command format to update the time zone setting:

```
chtz=local_timezone,date/time,date/time
```

For example, if you lived in Tucson, Arizona, where Mountain Standard Time is used, you would issue the following command to change to the POSIX format:

```
chtz MST7MDT,M3.2.0/2:00:00,M11.1.0/2:00:00
```

6. Add an entry in the .profile of the instance user so that the following environment is set:

```
export MALLOCOPTIONS=multiheap:16
```

Tip: If the instance user is not available, complete this step later, when the instance user becomes available.

7. Set the system to create full application core files. Issue the following command:

```
chdev -l sys0 -a fullcore=true -P
```

8. For communications with the server and Operations Center, make sure that the following ports are open on any firewalls that might exist:

- o For communications with the server, open port 1500.
- o For secure communications with the Operations Center, open port 11090 on the hub server.

If you are not using the default port values, make sure that the ports that you are using are open.

9. Enable TCP high-performance enhancements. Issue the following command:

```
no -p -o rfc1323=1
```

10. For optimal throughput and reliability, bond four 10 Gb Ethernet ports together. Use the System Management Interface Tool (SMIT) to bond the ports together by using Etherchannel. The following settings were used during testing:

```
mode          8023ad
auto_recovery yes          Enable automatic recovery after failover
backup_adapter NONE       Adapter used when whole channel fails
hash_mode     src_dst_port Determines how outgoing adapter is chosen
interval      long        Determines interval value for IEEE
                        802.3ad mode
mode          8023ad      EtherChannel mode of operation
netaddr       0           Address to ping
no_loss_failover yes      Enable lossless failover after ping
                        failure
num_retries   3           Times to retry ping before failing
retry_time    1           Wait time (in seconds) between pings
use_alt_addr  no          Enable Alternate EtherChannel Address
use_jumbo_frame no       Enable Gigabit Ethernet Jumbo Frames
```

11. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 1. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 1. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
-----------------	---------	-------	------------------------

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	FSIZE	Unlimited	ulimit -Hf
Maximum number of open files	nfile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Linux systems

Complete the following steps to install Linux x86_64 on the server system.

Before you begin

The operating system will be installed on the internal hard disks. Configure the internal hard disks by using a hardware RAID 1 array. For example, if you are configuring a small system, the two 300 GB internal disks are mirrored in RAID 1 so that a single 300 GB disk appears available to the operating system installer.

Procedure

1. Install Red Hat Enterprise Linux Version 7.1 or later, according to the manufacturer instructions. Obtain a bootable DVD that contains Red Hat Enterprise Linux Version 7.1 and start your system from this DVD. See the following guidance for installation options. If an item is not mentioned in the following list, leave the default selection.
 - a. After you start the DVD, choose Install or upgrade an existing system from the menu.
 - b. On the Welcome screen, select Test this media & install Red Hat Enterprise Linux 7.1.
 - c. Select your language and keyboard preferences.
 - d. Select your location to set the correct time zone.
 - e. Select Software Selection and then on the next screen, select Server with GUI.
 - f. From the installation summary page, click Installation Destination and verify the following items:
 - The local 300 GB disk is selected as the installation target.
 - Under Other Storage Options, Automatically configure partitioning is selected.
 - g. Click Done.
 - g. Click Begin Installation. After the installation starts, set the root password for your root user account.

After the installation is completed, restart the system and log in as the root user. Issue the `df` command to verify your basic partitioning. For example, on a test system, the initial partitioning produced the following result:

```
[root@tvapp02]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  50G  3.0G  48G   6% /
devtmpfs        32G   0    32G   0% /dev
tmpfs           32G  92K   32G   1% /dev/shm
tmpfs           32G  8.8M  32G   1% /run
tmpfs           32G   0    32G   0% /sys/fs/cgroup
/dev/mapper/rhel-home 220G  37M  220G   1% /home
/dev/sda1       497M 124M  373M  25% /boot
```

2. Configure your TCP/IP settings according to the operating system installation instructions.

For optimal throughput and reliability, consider bonding multiple network ports together. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection, which aggregates several subordinate ports into a single logical connection. The preferred method is to use a bond mode of 802.3ad, miimon setting of 100, and a `xmit_hash_policy` setting of layer3+4.

Restriction: To use an LACP network connection, you must have a network switch that supports LACP.

For additional instructions about configuring bonded network connections with Red Hat Enterprise Linux Version 7, see [Create a Channel Bonding Interface](#).

3. Open the `/etc/hosts` file and complete the following actions:

- o Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7 server.yourdomain.com server
```

- o Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1 localhost
```

4. Install components that are required for the server installation. Complete the following steps to create a Yellowdog Updater Modified (YUM) repository and install the prerequisite packages.

- a. Mount your Red Hat Enterprise Linux installation DVD to a system directory. For example, to mount it to the `/mnt` directory, issue the following command:

```
mount -t iso9660 -o ro /dev/cdrom /mnt
```

- b. Verify that the DVD mounted by issuing the mount command. You should see output similar to the following example:

```
/dev/sr0 on /mnt type iso9660
```

- c. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

If the `repos.d` directory does not exist, create it.

- d. List directory contents:

```
ls rhel-source.repo
```

- e. Rename the original repo file by issuing the `mv` command. For example:

```
mv rhel-source.repo rhel-source.repo.orig
```

- f. Create a new repo file by using a text editor. For example, to use the `vi` editor, issue the following command:

```
vi rhel71_dvd.repo
```

- g. Add the following lines to the new repo file. The `baseurl` parameter specifies your directory mount point:

```
[rhel71_dvd]
name=DVD Redhat Enterprise Linux 7.1
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

- h. Install the prerequisite package `ksh.x86_64`, by issuing the `yum` command. For example:

```
yum install ksh.x86_64
```

Exception: You do not need to install the `compat-libstdc++-33-3.2.3-69.el6.i686` and `libstdc++.i686` libraries for Red Hat Enterprise Linux Version 7.1.

5. When the software installation is complete, you can restore the original YUM repository values by completing the following steps:

- a. Unmount the Red Hat Enterprise Linux installation DVD by issuing the following command:

```
umount /mnt
```

- b. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

- c. Rename the repo file that you created:

```
mv rhel71_dvd.repo rhel71_dvd.repo.orig
```

- d. Rename the original file to the original name:

```
mv rhel-source.repo.orig rhel-source.repo
```

6. Determine whether kernel parameter changes are required. Complete the following steps:
 - a. Use the `sysctl -a` command to list the parameter values.
 - b. Analyze the results by using the guidelines in Table 1 to determine whether any changes are required.
 - c. If changes are required, set the parameters in the `/etc/sysctl.conf` file. The file changes are applied when the system is started.

Tip: Automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings. On Linux, the DB2® database software automatically adjusts interprocess communication (IPC) kernel parameter values to the preferred settings. For more information about kernel parameter settings, search for Linux kernel parameters in the IBM DB2 Version 11.1 product documentation.

Table 1. Linux kernel parameter optimum settings

Parameter	Description
<code>kernel.shmni</code>	The maximum number of segments.
<code>kernel.shmmax</code>	The maximum size of a shared memory segment (bytes). This parameter must be set before automatically starting the IBM Spectrum Protect™ server on system startup.
<code>kernel.shmall</code>	The maximum allocation of shared memory pages (pages).
<code>kernel.sem</code>	(SEMMSL) The maximum semaphores per array.
There are four values for the <code>kernel.sem</code> parameter.	(SEMMNS) The maximum semaphores per system.
	(SEMOPM) The maximum operations per semaphore call.
	(SEMMNI) The maximum number of arrays.
<code>kernel.msgmni</code>	The maximum number of system-wide message queues.
<code>kernel.msgmax</code>	The maximum size of messages (bytes).
<code>kernel.msgmnb</code>	The default maximum size of queue (bytes).
<code>kernel.randomize_va_space</code>	The <code>kernel.randomize_va_space</code> parameter configures the use of memory ASLR for the kernel. Disable ASLR because it can cause errors for the DB2 software. To learn more details about the Linux ASLR and DB2, see technote 1365583.
<code>vm.swappiness</code>	The <code>vm.swappiness</code> parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the DB2 product information.
<code>vm.overcommit_memory</code>	The <code>vm.overcommit_memory</code> parameter influences how much virtual memory the kernel permits allocating. For more information about kernel parameters, see the DB2 product information.

7. Open firewall ports to communicate with the server. Complete the following steps:
 - a. Determine the zone that is used by the network interface. The zone is public, by default.
Issue the following command:

```
# firewall-cmd --get-active-zones
public
  interfaces: ens4f0
```

- b. To use the default port address for communications with the server, open TCP/IP port 1500 in the Linux firewall.
Issue the following command:

```
firewall-cmd --zone=public --add-port=1500/tcp --permanent
```

If you want to use a value other than the default, you can specify a number in the range 1024 - 32767. If you open a port other than the default, you will need to specify that port when you run the configuration script.

- c. If you plan to use this system as a hub, open port 11090, which is the default port for secure (https) communications.

Issue the following command:

```
firewall-cmd --zone=public --add-port=11090/tcp --permanent
```

- d. Reload the firewall definitions for the changes to take effect.

Issue the following command:

```
firewall-cmd --reload
```

8. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 2. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 2. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf
Maximum number of open files	nofile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Windows systems

Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect™ server.

Procedure

1. Install Windows Server 2016 Standard Edition, according to the manufacturer instructions.
2. Change the Windows account control policies by completing the following steps.
 - a. Open the Local Security Policy editor by running secpol.msc.
 - b. Click Local Policies > Security Options and ensure that the following User Account Control policies are disabled:
 - Admin Approval Mode for the Built-in Administrator account
 - Run all administrators in Admin Approval Mode
3. Configure your TCP/IP settings according to installation instructions for the operating system.
4. Apply Windows updates and enable optional features by completing the following steps:
 - a. Apply the latest Windows Server 2016 updates.
 - b. Install and enable the Windows 2012 R2 feature Microsoft .NET Framework 3.5 from the Windows Server Manager.
 - c. If required, update the FC and Ethernet HBA device drivers to newer levels.
 - d. Install the multipath I/O driver that is appropriate for the disk system that you are using.

5. Open the default TCP/IP port, 1500, for communications with the IBM Spectrum Protect server. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Backup server port 1500" dir=in action=allow protocol=TCP localport=1500
```

6. On the Operations Center hub server, open the default port for secure (https) communications with the Operations Center. The port number is 11090. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Operations Center port 11090"  
dir=in action=allow protocol=TCP localport=11090
```

Configuring multipath I/O

You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.

- AIX systems
- Linux systems
- Windows systems

AIX systems

Procedure

1. Determine the Fibre Channel port address that you must use for the host definition on the disk subsystem. Issue the `lscfg` command for every port.

- On small and medium systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"  
lscfg -vps -l fcs1 | grep "Network Address"
```

- On large systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"  
lscfg -vps -l fcs1 | grep "Network Address"  
lscfg -vps -l fcs2 | grep "Network Address"  
lscfg -vps -l fcs3 | grep "Network Address"
```

2. Ensure that the following AIX® file sets are installed:

- `devices.common.IBM.mpio.rte`
- `devices.fcp.disk.array.rte`
- `devices.fcp.disk.rte`

3. Issue the `cfgmgr` command to have AIX rescan the hardware and discover available disks. For example:

```
cfgmgr
```

4. To list the available disks, issue the following command:

```
lsdev -Ccdisk
```

You should see output similar to the following:

```
hdisk0 Available 00-00-00 SAS Disk Drive  
hdisk1 Available 00-00-00 SAS Disk Drive  
hdisk2 Available 01-00-00 SAS Disk Drive  
hdisk3 Available 01-00-00 SAS Disk Drive  
hdisk4 Available 06-01-02 MPIO IBM 2076 FC Disk  
hdisk5 Available 07-01-02 MPIO IBM 2076 FC Disk  
...
```

5. Use the output from the `lsdev` command to identify and list device IDs for each disk device.

For example, a device ID could be `hdisk4`. Save the list of device IDs to use when you create file systems for the IBM Spectrum Protect™ server.

6. Correlate the SCSI device IDs to specific disk LUNs from the disk system by listing detailed information about all physical volumes in the system. Issue the following command:

```
lspv -u
```

On an IBM® Storwize® system, the following information is an example of what is shown for each device:

```
hdisk4 00f8cf083fd97327 None active  
33213600507630081010578000000000003004214503IBMfcp
```

In the example, 6005076300810105780000000000030 is the UID for the volume, as reported by the Storwize management interface.

To verify disk size in megabytes and compare the value with what is listed for the system, issue the following command:

```
bootinfo -s hdisk4
```

Linux systems

Procedure

1. Edit the `/etc/multipath.conf` file to enable multipathing for Linux hosts. If the `multipath.conf` file does not exist, you can create it by issuing the following command:

```
multipathconf --enable
```

The following parameters were set in `multipath.conf` for testing on an IBM Storwize® system:

```
defaults {
    user_friendly_names no
}

devices {
    device {
        vendor "IBM "
        product "2145"
        path_grouping_policy group_by_prio
        user_friendly_names no
        path_selector "round-robin 0"
        prio "alua"
        path_checker "tur"
        failback "immediate"
        no_path_retry 5
        rr_weight uniform
        rr_min_io_rq "1"
        dev_loss_tmo 120
    }
}
```

2. Set the multipath option to start when the system is started. Issue the following commands:

```
systemctl enable multipathd.service
systemctl start multipathd.service
```

3. To verify that disks are visible to the operating system and are managed by multipath, issue the following command:

```
multipath -l
```

4. Ensure that each device is listed and that it has as many paths as you expect. You can use size and device ID information to identify which disks are listed.

For example, the following output shows that a 2 TB disk has two path groups and four active paths. The 2 TB size confirms that the disk corresponds to a pool file system. Use part of the long device ID number (12, in this example) to search for the volume on the disk-system management interface.

```
[root@tapsrv01 code]# multipath -l
36005076802810c50980000000000012 dm-43 IBM,2145
 size=2.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='round-robin 0' prio=0 status=active
|  |- 2:0:1:18 sdcw 70:64 active undef running
|  `-- 4:0:0:18 sdgb 131:112 active undef running
`--+- policy='round-robin 0' prio=0 status=enabled
    |- 1:0:1:18 sdat 66:208 active undef running
    `-- 3:0:0:18 sddy 128:0 active undef running
```

- a. If needed, correct disk LUN host assignments and force a bus rescan. For example:

```
echo "- - -" > /sys/class/scsi_host/host0/scan
echo "- - -" > /sys/class/scsi_host/host1/scan
echo "- - -" > /sys/class/scsi_host/host2/scan
```

You can also restart the system to rescan disk LUN host assignments.

- b. Confirm that disks are now available for multipath I/O by reissuing the `multipath -l` command.
5. Use the multipath output to identify and list device IDs for each disk device.

For example, the device ID for your 2 TB disk is `36005076802810c509800000000000012`.

Save the list of device IDs to use in the next step.

Windows systems

Procedure

1. Ensure that the Multipath I/O feature is installed. If needed, install additional vendor-specific multipath drivers.
2. To verify that disks are visible to the operating system and are managed by multipath I/O, issue the following command:

```
c:\program files\IBM\SDDDSM\datapath.exe query device
```

3. Review the multipath output and ensure that each device is listed and that it has as many paths as you expect. You can use size and device serial information to identify which disks are listed.

For example, by using part of the long device serial number (34, in this example) you can search for the volume on the disk-system management interface. The 2 TB size confirms that the disk corresponds to a storage pool file system.

```
DEV#: 4 DEVICE NAME: Disk5 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 600507630081010578000000000000034 LUN SIZE: 2.0TB
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 0 0
1 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 27176 0
2 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 28494 0
3 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 0 0
```

4. Create a list of disk device IDs by using the serial numbers that are returned from the multipath output in the previous step.

For example, the device ID for your 2 TB disk is `600507630081010578000000000000034`

Save the list of device IDs to use in the next step.

5. To bring new disks online and clear the read-only attribute, run `diskpart.exe` with the following commands. Repeat for each of the disks:

```
diskpart
select Disk 1
online disk
attribute disk clear readonly
select Disk 2
online disk
attribute disk clear readonly
< ... >
select Disk 49
online disk
attribute disk clear readonly
exit
```

Creating the user ID for the server

Create the user ID that owns the IBM Spectrum Protect™ server instance. You specify this user ID when you create the server instance during initial configuration of the server.

About this task

You can specify only lowercase letters (a-z), numerals (0-9), and the underscore character (`_`) for the user ID. The user ID and group name must comply with the following rules:

- The length must be 8 characters or fewer.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

Procedure

1. Use operating system commands to create a user ID.

- o **AIX** | **Linux** Create a group and user ID in the home directory of the user that owns the server instance.

For example, to create the user ID `tsminst1` in group `tsmsrvrs` with a password of `tsminst1`, issue the following commands from an administrative user ID:

AIX

```
mkgroup id=1001 tsmsrvrs
mkuser id=1002 pgrp=tsmsrvrs home=/home/tsminst1 tsminst1
passwd tsminst1
```

Linux

```
groupadd tsmsrvrs
useradd -d /home/tsminst1 -m -g tsmsrvrs -s /bin/bash tsminst1
passwd tsminst1
```

Log off, and then log in to your system. Change to the user account that you created. Use an interactive login program, such as `telnet`, so that you are prompted for the password and can change it if necessary.

- o **Windows** Create a user ID and then add the new ID to the Administrators group. For example, to create the user ID `tsminst1`, issue the following command:

```
net user tsminst1 * /add
```

After you create and verify a password for the new user, add the user ID to the Administrators group by issuing the following commands:

```
net localgroup Administrators tsminst1 /add
net localgroup DB2ADMNS tsminst1 /add
```

2. Log off the new user ID.

Preparing file systems for the server

You must complete file system configuration for the disk storage to be used by the server.

- Preparing file systems on AIX systems
You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.
- Preparing file systems on Linux systems
You must format `ext4` or `xfs` file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.
- Preparing file systems on Windows systems
You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.

Preparing file systems on AIX systems

You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.

Procedure

1. Increase the queue depth and maximum transfer size for all of the available `hdiskX` disks. Issue the following commands for each disk:

```
chdev -l hdisk4 -a max_transfer=0x100000
chdev -l hdisk4 -a queue_depth=32
chdev -l hdisk4 -a reserve_policy=no_reserve
chdev -l hdisk4 -a algorithm=round_robin
```

Do not run these commands for operating system internal disks, for example, `hdisk0`.

2. Create volume groups for the IBM Spectrum Protect™ database, active log, archive log, database backup, and storage pool. Issue the `mkvg` command, specifying the device IDs for corresponding disks that you previously identified.

For example, if the device names *hdisk4*, *hdisk5*, and *hdisk6* correspond to database disks, include them in the database volume group and so on.

System size: The following commands are based on the medium system configuration. For small and large systems, you must adjust the syntax as required.

```
mkvg -S -y tsmdb hdisk2 hdisk3 hdisk4
mkvg -S -y tsmactlog hdisk5
mkvg -S -y tsmarchlog hdisk6
mkvg -S -y tsmdbback hdisk7 hdisk8 hdisk9 hdisk10
mkvg -S -y tsmstgpool hdisk11 hdisk12 hdisk13 hdisk14 ... hdisk49
```

3. Determine the physical volume names and the number of free physical partitions to use when you create logical volumes. Issue the `lsvg` for each volume group that you created in the previous step.

For example:

```
lsvg -p tsmdb
```

The output is similar to the following. The *FREE PPs* column represents the free physical partitions:

```
tsmdb:
PV_NAME  PV STATE  TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk4   active    1631       1631      327..326..326..326..326
hdisk5   active    1631       1631      327..326..326..326..326
hdisk6   active    1631       1631      327..326..326..326..326
```

4. Create logical volumes in each volume group by using the `mklv` command. The volume size, volume group, and device names vary, depending on the size of your system and variations in your disk configuration.

For example, to create the volumes for the IBM Spectrum Protect database on a medium system, issue the following commands:

```
mklv -y tsmdb00 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk2
mklv -y tsmdb01 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk3
mklv -y tsmdb02 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk4
```

5. Format file systems in each logical volume by using the `crfs` command.

For example, to format file systems for the database on a medium system, issue the following commands:

```
crfs -v jfs2 -d tsmdb00 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace00 -A yes
crfs -v jfs2 -d tsmdb01 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace01 -A yes
crfs -v jfs2 -d tsmdb02 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace02 -A yes
```

6. Mount all of the newly created file systems by issuing the following command:

```
mount -a
```

7. List all file systems by issuing the `df` command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example of command output shows that the amount of used space is typically 1%:

```
tapsrv07> df -g /tsminst1/*
Filesystem      GB blocks  Free    %Used  Iused  %Iused  Mounted on
/dev/tsmact00   195.12    194.59  1%      4      1%      /tsminst1/TSMalog
```

8. Verify that the user ID you created in *Creating the user ID for the server* has read and write access to the directories for the IBM Spectrum Protect server.

Preparing file systems on Linux systems

You must format ext4 or xfs file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Using the list of device IDs that you generated previously, issue the `mkfs` command to create and format a file system for each storage LUN device. Specify the device ID in the command. See the following examples. For the database, format ext4 file systems:

```
mkfs -t ext4 -T largefile -m 2 /dev/mapper/36005076802810c50980000000000012
```

For storage pool LUNs, format xfs file systems:

```
mkfs -t xfs /dev/mapper/3600507630081010578000000000002c3
```

You might issue the mkfs command as many as 50 times, depending on how many different devices you have.

2. Create mount point directories for file systems.

Issue the mkdir command for each directory that you must create. Use the directory values that you recorded in the planning worksheets.

For example, to create the server instance directory by using the default value, issue the following command:

```
mkdir /tsminst1
```

Repeat the mkdir command for each file system.

3. Add an entry in the /etc/fstab file for each file system so that file systems are mounted automatically when the server is started.

For example:

```
/dev/mapper/36005076802810c50980000000000012 /tsminst1/TSMdbspace00 ext4 defaults 0 0
```

4. Mount the file systems that you added to the /etc/fstab file by issuing the mount -a command.

5. List all file systems by issuing the df command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example on an IBM® Storwize® system shows that the amount of used space is typically 1%:

```
[root@tapsrv04 ~]# df -h /tsminst1/*
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/36005076300810105780000000000003 134G  188M 132G   1% /tsminst1/TSMalog
```

6. Verify that the user ID you created in Creating the user ID for the server has read and write access to the directories for IBM Spectrum Protect.

Preparing file systems on Windows systems

You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Create mount point directories for file systems.

Issue the md command for each directory that you must create. Use the directory values that you recorded in the planning worksheets. For example, to create the server instance directory by using the default value, issue the following command:

```
md c:\tsminst1
```

Repeat the md command for each file system.

2. Create a volume for every disk LUN that is mapped to a directory under the server instance directory by using the Windows volume manager.

Go to Server Manager > File and Storage Services and complete the following steps for each disk that corresponds to the LUN mapping that was created in the previous step:

a. Bring the disk online.

b. Initialize the disk to the GPT basic type, which is the default.

c. Create a simple volume that occupies all of the space on the disk. Format the file system by using NTFS, and assign a label that matches the purpose of the volume, such as TSMfile00. Do not assign the new volume to a drive letter. Instead, map the volume to a directory under the instance directory, such as C:\tsminst1\TSMfile00.

Tip: Determine the volume label and directory mapping labels based on the size of the disk that is reported.

3. Verify that file systems are mounted at the correct LUN and correct mount point. List all file systems by issuing the mountvol command and then review the output. For example:

```
\\?\Volume{8ffb9678-3216-474c-a021-20e420816a92}\
C:\tsminst1\TSMdbspace00\
```

4. After the disk configuration is complete, restart the system.

What to do next

You can confirm the amount of free space for each volume by using Windows Explorer.

Installing the server and Operations Center

Use the IBM® Installation Manager graphical wizard to install the components.

- Installing on AIX and Linux systems
Install the IBM Spectrum Protect™ server and the Operations Center on the first server system.
- Installing on Windows systems
Install the IBM Spectrum Protect server and the Operations Center on the first server system.

Installing on AIX® and Linux systems

Install the IBM Spectrum Protect™ server and the Operations Center on the first server system.

Before you begin

Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

1. **AIX** Verify that the required RPM files are installed on your system.

See Installing prerequisite RPM files for the graphical wizard for details.

2. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042992.
3. Go to Passport Advantage® and download the package file to an empty directory of your choice.
4. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

5. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file.

6. **AIX** Ensure that the following command is enabled so that the wizards work properly:

```
lsuser
```

By default, the command is enabled.

7. Change to the directory where you placed the executable file.
8. Start the installation wizard by issuing the following command:

```
./install.sh
```

When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.

- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.
- Installing prerequisite RPM files for the graphical wizard
RPM files are required for the IBM Installation Manager graphical wizard.

Related tasks:

- 🔗 Other methods for installing IBM Spectrum Protect components (AIX)
- 🔗 Other methods for installing IBM Spectrum Protect components (Linux)

Installing on Windows systems

Install the IBM Spectrum Protect™ server and the Operations Center on the first server system.

Before you begin

Make sure that the following prerequisites are met:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

1. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042993.
2. Go to Passport Advantage® and download the package file to an empty directory of your choice.
3. Change to the directory where you placed the executable file.
4. Double-click the executable file to extract to the current directory.
5. In the directory where the installation files were extracted, start the installation wizard by double-clicking the install.bat file. When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.

- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.

Related tasks:

- 🔗 Other methods for installing IBM Spectrum Protect components

Configuring the server and the Operations Center

After you install the components, complete the configuration for the IBM Spectrum Protect™ server and the Operations Center.

- Configuring the server instance
Use the IBM Spectrum Protect server instance configuration wizard to complete the initial configuration of the server.
- Installing the backup-archive client
As a best practice, install the IBM Spectrum Protect backup-archive client on the server system so that the administrative command-line client and scheduler are available.
- Setting options for the server
Review the server options file that is installed with the IBM Spectrum Protect server to verify that the correct values are set for your system.
- Configuring secure communications with Transport Layer Security
To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security

(TLS) is enabled on the IBM Spectrum Protect server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.

- **Configuring the Operations Center**
After you install the Operations Center, complete the following configuration steps to start managing your storage environment.
- **Registering the product license**
To register your license for the IBM Spectrum Protect product, use the REGISTER LICENSE command.
- **Configuring data deduplication**
Create a directory-container storage pool and at least one directory to use inline data deduplication.
- **Defining data retention rules for your business**
After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.
- **Defining schedules for server maintenance activities**
Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations Center command builder.
- **Defining client schedules**
Use the Operations Center to create schedules for client operations.

Configuring the server instance

Use the IBM Spectrum Protect™ server instance configuration wizard to complete the initial configuration of the server.

Before you begin

Ensure that the following requirements are met:

AIX | **Linux**

- The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/directory`. Also, ensure that the SSH daemon service has access rights to connect to the system by using the `localhost` value.
- You must be able to log in to IBM Spectrum Protect with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

Windows

Verify that the remote registry service is started by completing the following steps:

1. Click Start > Administrative Tools > Services. In the Services window, select Remote Registry. If it is not started, click Start.
2. Ensure that port 137, 139, and 445 are not blocked by a firewall:
 - a. Click Start > Control Panel > Windows Firewall.
 - b. Select Advanced Settings.
 - c. Select Inbound Rules.
 - d. Select New Rule.
 - e. Create a port rule for TCP ports 137, 139, and 445 to allow connections for domain and private networks.
3. Configure the user account control by accessing the local security policy options and completing the following steps.
 - a. Click Start > Administrative Tools > Local Security Policy. Expand Local Policies > Security Options.
 - b. If not already enabled, enable the built-in administrator account by selecting Accounts: Administrator account status > Enable > OK.
 - c. If not already disabled, disable user account control for all Windows administrators by selecting User Account Control: Run all administrators in Admin Approval Mode > Disable > OK.
 - d. If not already disabled, disable the User Account Control for the built-in Administrator account by selecting User Account Control: Admin Approval Mode for the Built-in Administrator Account > Disable > OK.
4. If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

About this task

The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Procedure

1. Start the local version of the wizard.
 - o **AIX** | **Linux** Open the ds micfgx program in the /opt/tivoli/tsm/server/bin directory. This wizard can be only run as a root user.
 - o **Windows** Click Start > All Programs > IBM Spectrum Protect > Configuration Wizard.
2. Follow the instructions to complete the configuration. Use the information that you recorded in Planning worksheets during IBM Spectrum Protect system set up to specify directories and options in the wizard.

AIX | **Linux** On the Server Information window, set the server to start automatically by using the instance user ID when the system boots.

Windows By using the configuration wizard, the server is set to start automatically when rebooted.

Installing the backup-archive client

As a best practice, install the IBM Spectrum Protect™ backup-archive client on the server system so that the administrative command-line client and scheduler are available.

Procedure

To install the backup-archive client, follow the installation instructions for your operating system.

- Install UNIX and Linux backup-archive clients
- Installing the Windows client for the first time

Setting options for the server

Review the server options file that is installed with the IBM Spectrum Protect™ server to verify that the correct values are set for your system.

Procedure

1. Go to the server instance directory and open the dsmserv.opt file.
2. Review the values in the following table and verify your server option settings, based on system size.

Server option	Small system value	Medium system value	Large system value
ACTIVELOGDIRECTORY	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
ACTIVELOGSIZE	131072	131072	262144
ARCHLOGCOMPRESS	Yes	No	No
ARCHLOGDIRECTORY	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
COMMMETHOD	TCP/IP	TCP/IP	TCP/IP
COMMTIMEOUT	3600	3600	3600
DEDUPREQUIRESBACKUP	No	No	No
DEVCONFIG	devconf.dat	devconf.dat	devconf.dat
EXPINTERVAL	0	0	0
IDLETIMEOUT	60	60	60
MAXSESSIONS	250	500	1000
NUMOPENVOLSALLOWED	20	20	20
TCPADMINPORT	1500	1500	1500
TCPPORT	1500	1500	1500
VOLUMEHISTORY	volhist.dat	volhist.dat	volhist.dat

Update server option settings if necessary, to match the values in the table. To make updates, close the dsmserv.opt file and use the SETOPT command from the administrative command-line interface to set the options.

For example, to update the IDLETIMEOUT option to 60, issue the following command:

```
setopt idletimeout 60
```

3. To configure secure communications for the server, clients, and the Operations Center, verify the options in the following table.

Server option	All system sizes
SSLFIPSMODE	NO
TCPPOINT	Specify the port number on which the server waits for requests for TCP/IP and SSL-enabled sessions from the client.
TCPADMINPORT	Specify the port address on which the server waits for requests for TCP/IP and SSL-enabled sessions from the command-line administrative client.

If any of the option values must be updated, edit the dsmserv.opt file by using the following guidelines:

- o Remove the asterisk at the beginning of a line to enable an option.
- o On each line, enter only one option and the specified value for the option.
- o If an option occurs in multiple entries in the file, the server uses the last entry.

Save your changes and close the file. If you edit the dsmserv.opt file directly, you must restart the server for the changes to take effect.

Related reference:

[Server options reference](#)

[SETOPT \(Set a server option for dynamic update\)](#)

Configuring secure communications with Transport Layer Security

To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) is enabled on the IBM Spectrum Protect™ server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.

About this task

Beginning with IBM Spectrum Protect Version 8.1.2, SSL is enabled by default, and the IBM Spectrum Protect server and backup-archive client are automatically configured to communicate with each other by using the TLS 1.2 protocol.

As shown in the following figure, you can manually configure secure communications between the server and backup-archive client by setting options in the server and client options files, and then transferring the self-signed certificate that is generated on the server to the client. Alternatively, you can obtain and transfer a unique certificate that is signed by a certificate authority (CA).



For more information about configuring the server and clients for SSL or TLS communications, see [Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL](#).

Configuring the Operations Center

After you install the Operations Center, complete the following configuration steps to start managing your storage environment.

Before you begin

When you connect to the Operations Center for the first time, you must provide the following information:

- Connection information for the server that you want to designate as a hub server
- Login credentials for an administrator ID that is defined for that server

Procedure

1. Designate the hub server. In a web browser, enter the following address:

```
https://hostname:secure_port/oc
```

where:

- *hostname* represents the name of the computer where the Operations Center is installed
- *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer

For example, if your host name is `tsm.storage.mylocation.com` and you are using the default secure port for the Operations Center, which is 11090, the address is:

```
https://tsm.storage.mylocation.com:11090/oc
```

When you log in to the Operations Center for the first time, a wizard guides you through an initial configuration to set up a new administrator with system authority on the server.

2. Set up secure communications between the Operations Center and the hub server by configuring the Secure Sockets Layer (SSL) protocol.

Follow the instructions in [Securing communications between the Operations Center and the hub server](#).

3. Optional: To receive a daily email report that summarizes system status, configure your email settings in the Operations Center.

Follow the instructions in [Tracking system status by using email reports](#).

- Securing communications between the Operations Center and the hub server
To secure communications between the Operations Center and the hub server, add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.

Registering the product license


To register your license for the IBM Spectrum Protect™ product, use the REGISTER LICENSE command.

About this task

Licenses are stored in enrollment certificate files, which contain licensing information for the product. The enrollment certificate files are on the installation media, and are placed on the server during installation. When you register the product, the licenses are stored in a NODELOCK file within the current directory.

Procedure

Register a license by specifying the name of the enrollment certificate file that contains the license. To use the Operations Center command builder for this task, complete the following steps.

1. Open the Operations Center.
2. Open the Operations Center command builder by hovering over the settings icon  and clicking Command Builder.
3. Issue the REGISTER LICENSE command. For example, to register a base IBM Spectrum Protect license, issue the following command:

```
register license file=tsmbasic.lic
```

What to do next

Save the installation media that contains your enrollment certificate files. You might need to register your license again if, for example, one of the following conditions occur:

- The server is moved to a different computer.

- The NODELOCK file is corrupted. The server stores license information in the NODELOCK file, which is in the directory from which the server is started.
- **Linux** If you change the processor chip that is associated with the server on which the server is installed.

Related reference:

[REGISTER LICENSE](#) (Register a new license)

Configuring data deduplication

Create a directory-container storage pool and at least one directory to use inline data deduplication.

Before you begin

Use the storage pool directory information that you recorded in Planning worksheets for this task.

Procedure

1. Open the Operations Center.
2. On the Operations Center menu bar, hover over Storage.
3. From the list that is displayed, click Storage Pools.
4. Click the +Storage Pool button.
5. Complete the steps in the Add Storage Pool wizard:
 - To use inline data deduplication, select a Directory storage pool under Container-based storage.
 - When you configure directories for the directory-container storage pool, specify the directory paths that you created for storage during system setup.
6. After you configure the new directory-container storage pool, click Close & View Policies to update a management class and start using the storage pool.

Defining data retention rules for your business

After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.

Procedure

1. On the Services page of the Operations Center, select the STANDARD domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab. The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.
3. Click the Configure toggle, and make the following changes:
 - Change the backup destination for the STANDARD management class to the directory-container storage pool.
 - Change the value for the Backups column to No limit.
 - Change the retention period. Set the Keep Extra Backups column to 30 days or more, depending on your business requirements.
4. Save your changes and click the Configure toggle again so that the policy set is no longer editable.
5. Activate the policy set by clicking Activate.

Related tasks:

Specifying rules for backing up and archiving client data

Defining schedules for server maintenance activities

Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations Center command builder.

About this task

Schedule server maintenance operations to run after client backup operations. You can control the timing of schedules by setting the start time in combination with the duration time for each operation.

The following example shows how you can schedule server maintenance processes in combination with the client backup schedule for a multisite disk solution.

Operation	Schedule
Client backup	Starts at 22:00.
Node replication	Starts at 08:00, or 10 hours after the beginning of the client backup.
Processing for database and disaster recovery files	<ul style="list-style-type: none"> Database backup starts at 11:00, or 13 hours after the beginning of the client backup. This process runs until completion. Device configuration information and volume history backup starts at 17:00, or 6 hours after the start of the database backup. Volume history deletion starts at 20:00, or 9 hours after the start of the database backup.
Inventory expiration	Starts at 12:00, or 14 hours after the beginning of the client backup window. This process runs until completion.

Procedure

After you configure the device class for the database backup operations, create schedules for database backup and other required maintenance operations by using the DEFINE SCHEDULE command. Depending on the size of your environment, you might need to adjust the start times for each schedule in the example.

1. Define a device class for the backup operations. For example, use the DEFINE DEVCLASS command to create a device class that is named DBBACK_FILEDEV:

```
define devclass dbback_filedev devtype=file
    directory=db_backup_directories
```

where *db_backup_directories* is a list of the directories that you created for the database backup.

AIX **Linux** For example, if you have four directories for database backups, starting with /tsminst1/TSMbkup00, issue the following command:

```
define devclass dbback_filedev devtype=file
    directory=/tsminst1/TSMbkup00,
    /tsminst1/TSMbkup01,/tsminst1/TSMbkup02,
    /tsminst1/TSMbkup03"
```

Windows For example, if you have four directories for database backups, starting with C:\tsminst1\TSMbkup00, issue the following command:

```
define devclass dbback_filedev devtype=file
    directory="c:\tsminst1\TSMbkup00,
    c:\tsminst1\TSMbkup01,c:\tsminst1\TSMbkup02,
    c:\tsminst1\TSMbkup03"
```

2. Set the device class for automatic database backup operations. Use the SET DBRECOVERY command to specify the device class that you created in the preceding step. For example, if the device class is dbback_filedev, issue the following command:

```
set dbrecovery dbback_filedev
```

3. Create schedules for the maintenance operations by using the DEFINE SCHEDULE command. See the following table for the required operations with examples of the commands.

Tip: You create the schedule for replication separately in a later step, when you use the Operations Center to configure replication.

Operation	Example command
-----------	-----------------

Operation	Example command
Back up the database.	<p>Create a schedule to run the BACKUP DB command. If you are configuring a small system, set the COMPRESS parameter to YES. For example, on a small system, issue the following command to create a backup schedule that uses the new device class:</p> <pre>define schedule DBBACKUP type=admin cmd="backup db devclass=dbback_filedev type=full numstreams=3 wait=yes compress=yes" active=yes desc="Back up the database." startdate=today starttime=11:00:00 duration=45 durunits=minutes</pre>
Back up the device configuration information.	<p>Create a schedule to run the BACKUP DEVCONFIG command:</p> <pre>define schedule DEVCONFIGBKUP type=admin cmd="backup devconfig filenames=devconfig.dat" active=yes desc="Backup the device configuration file." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>
Back up the volume history.	<p>Create a schedule to run the BACKUP VOLHISTORY command:</p> <pre>define schedule VOLHISTBKUP type=admin cmd="backup volhistory filenames=volhist.dat" active=yes desc="Back up the volume history." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>
Remove older versions of database backups that are no longer required.	<p>Create a schedule to run the DELETE VOLHISTORY command:</p> <pre>define schedule DELVOLHIST type=admin cmd="delete volhistory type=dbb todate=today-6 totime=now" active=yes desc="Remove old database backups." startdate=today starttime=20:00:00 duration=45 durunits=minutes</pre>
Remove objects that exceed their allowed retention.	<p>Create a schedule to run the EXPIRE INVENTORY command. Set the RESOURCE parameter based on the system size that you are configuring:</p> <ul style="list-style-type: none"> o Small systems: 10 o Medium systems: 30 o Large systems: 40 <p>For example, on a medium-sized system, issue the following command to create a schedule that is named EXPINVENTORY:</p> <pre>define schedule EXPINVENTORY type=admin cmd="expire inventory wait=yes resource=30 duration=120" active=yes desc="Remove expired objects." startdate=today starttime=12:00:00 duration=45 durunits=minutes</pre>

What to do next

After you create schedules for the server maintenance tasks, you can view them in the Operations Center by completing the following steps:

1. On the Operations Center menu bar, hover over Servers.
2. Click Maintenance.

Related reference:

➔ DEFINE SCHEDULE (Define a schedule for an administrative command)

Defining client schedules

Use the Operations Center to create schedules for client operations.

Procedure

1. On the Operations Center menu bar, hover over Clients.
2. Click Schedules.
3. Click +Schedule.
4. Complete the steps in the Create Schedule wizard. Set client backup schedules to start at 22:00, based on the server maintenance activities that you scheduled in Defining schedules for server maintenance activities.

Installing and configuring backup-archive clients

Following the successful setup of your IBM Spectrum Protect™ server system, install and configure client software to begin backing up data.

Procedure

To install the backup-archive client, follow the installation instructions for your operating system.

- Install UNIX and Linux backup-archive clients
- Installing the Windows client for the first time

What to do next

Register and assign your clients to schedules.

- Registering and assigning clients to schedules
Add and register your clients through the Operations Center by using the Add Client wizard.
- Installing the client management service
Install the client management service for backup-archive clients that run on Linux and Windows operating systems. The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Registering and assigning clients to schedules

Add and register your clients through the Operations Center by using the Add Client wizard.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see technote 7048963.

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see Authenticating users by using an Active Directory database.

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the REGISTER NODE command and specify the USERID parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see Register a node.

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:
 - a. On the Operations Center menu bar, click Clients.
 - b. In the Clients table, click +Client.
 - c. Complete the steps in the Add Client wizard:
 - i. Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the Enable check box.
 - ii. In the Configuration window, copy the TCPSERVERADDRESS, TCPPORT, NODENAME, and DEDUPLICATION option values.
Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
 - iii. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
 - iv. Set how risks are displayed for the client by specifying the at-risk setting.
 - v. Click Add Client.

Installing the client management service

Install the client management service for backup-archive clients that run on Linux and Windows operating systems. The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Procedure

Install the client management service on the same computer as the backup-archive client by completing the following steps:

1. Download the installation package for the client management service from an IBM® download site such as IBM Passport Advantage® or IBM Fix Central. Look for a file name that is similar to `<version>-IBM_Spectrum_Protect-CMS-operating_system.bin`.
 2. Create a directory on the client system that you want to manage, and copy the installation package there.
 3. Extract the contents of the installation package file.
 4. Run the installation batch file from the directory where you extracted the installation and associated files. This is the directory that you created in step 2.
 5. To install the client management service, follow the instructions in the IBM Installation Manager wizard. If IBM Installation Manager is not already installed on the client system, you must select both IBM Installation Manager and IBM Spectrum Protect™ Client Management Services.
- Verifying that the client management service is installed correctly
Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.
 - Configuring the Operations Center to use the client management service
If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Related tasks:

[↗ Configuring the client management service for custom client installations](#)

Verifying that the client management service is installed correctly

Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.

Procedure

On the client system, at the command line, run the following commands to view the configuration of the client management service:

- On Linux client systems, issue the following command:

```
client_install_dir/cms/bin/CmsConfig.sh list
```

where `client_install_dir` is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
/opt/tivoli/tsm/cms/bin/CmsConfig.sh list
```

The output is similar to the following text:

Listing CMS configuration

```
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys

  Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

- On Windows client systems, issue the following command:

```
client_install_dir\cms\bin\CmsConfig.bat list
```

where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
C:"Program Files"\Tivoli\TSM\cms\bin\CmsConfig.bat list
```

The output is similar to the following text:

Listing CMS configuration

```
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt

  Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Log File: C:\Program Files\Tivoli\TSM\baclient\dsmsched.log
            en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

If the client management service is correctly installed and configured, the output displays the location of the error log file. The output text is extracted from the following configuration file:

- On Linux client systems:

```
client_install_dir/cms/Liberty/usr/servers/cmsServer/client-configuration.xml
```

- On Windows client systems:

```
client_install_dir\cms\Liberty\usr\servers\cmsServer\client-configuration.xml
```

If the output does not contain any entries, you must configure the *client-configuration.xml* file. For instructions to configure this file, see *Configuring the client management service for custom client installations*. You can use the *CmsConfig verify* command to verify that a node definition is correctly created in the *client-configuration.xml* file.

Configuring the Operations Center to use the client management service

If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Before you begin

Ensure that the client management service is installed and started on the client system. Verify whether the default configuration is used. The default configuration is not used if either of the following conditions is met:

- The client management service does not use the default port number, 9028.
- The backup-archive client is not accessed by the same IP address as the client system where the backup-archive client is installed. For example, a different IP address might be used in the following situations:
 - The computer system has two network cards. The backup-archive client is configured to communicate on one network, while the client management service communicates on the other network.

- o The client system is configured with the Dynamic Host Configuration Protocol (DHCP). As a result, the client system is dynamically assigned an IP address, which is saved on the server during the previous backup-archive client operation. When the client system is restarted, the client system might be assigned a different IP address. To ensure that the Operations Center can always find the client system, you specify a fully qualified domain name.

Procedure

To configure the Operations Center to use the client management service, complete the following steps:

1. On the Clients page of the Operations Center, select the client.
2. Click Details > Properties.
3. In the Remote diagnostics URL field in the General section, specify the URL for the client management service on the client system. The address must start with `https`. The following table shows examples of the remote diagnostics URL.

Type of URL	Example
With DNS host name and default port, 9028	<code>https://server.example.com</code>
With DNS host name and non-default port	<code>https://server.example.com:1599</code>
With IP address and non-default port	<code>https://192.0.2.0:1599</code>

4. Click Save.

What to do next

You can access client diagnostic information such as client log files from the Diagnosis tab in the Operations Center.

Configuring the second server

After you complete the configuration for the first server in your system, configure the second server.

Procedure

Complete the instructions in the following sections:

1. Configure a second server that is the same as the first server by completing the instructions in the following sections:
 - a. Setting up the system
 - b. Installing the server and Operations Center

Only one server in the multisite disk solution is configured as the hub server, so you do not need to install the Operations Center on the second server. When you select the installation packages to install on the second server, do not select the Operations Center.

- c. Configuring the server and the Operations Center

Skip the tasks for configuring the Operations Center.

- d. Installing and configuring backup-archive clients
2. Configuring SSL communications between the hub server and a spoke server
 3. Adding the second server as a spoke
 4. Enabling replication

Configuring SSL communications between the hub server and a spoke server

To secure communications between the hub server and a spoke server by using the Transport Layer Security (TLS) protocol, you must define the certificate of the spoke server to the hub server.

About this task

The hub server receives status and alert information from the spoke server and shows this information in the Operations Center. To receive the status and alert information from the spoke server, the certificate of the spoke server must be added to the truststore file of the hub server. You must also configure the Operations Center to monitor the spoke server.

To enable other functions of the Operations Center, such as the automatic deployment of client updates, the certificate of the hub server must be added to the truststore file of the spoke server.

Procedure

1. Complete the following steps to define the certificate of the spoke server to the hub server:
 - a. On the spoke server, change to the directory of the spoke server instance.
 - b. Specify the required cert256.arm certificate as the default certificate in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed  
-label "TSM Server SelfSigned SHA Key"
```

- c. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- d. Securely transfer the cert256.arm file of the spoke server to the hub server.
- e. On the hub server, change to the directory of the hub server instance.
- f. Define the spoke server certificate to the hub server. Issue the following command from the hub server instance directory, where *spoke_servername* is the name of the spoke server, and *spoke_cert256.arm* is the file name of the spoke server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label spoke_servername -file spoke_cert256.arm
```

2. Complete the following steps to define the certificate of the hub server to the spoke server:
 - a. On the hub server, change to the directory of the hub server instance.
 - b. Specify the required cert256.arm certificate as the default certificate in the key database file of the hub server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed  
-label "TSM Server SelfSigned SHA Key"
```

- c. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- d. Securely transfer the cert256.arm file of the hub server to the spoke server.
- e. On the spoke server, change to the directory of the spoke server instance.
- f. Define the hub server certificate to the spoke server. Issue the following command from the spoke server instance directory, where *hub_servername* is the name of the hub server, and *hub_cert256.arm* is the file name of the hub server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label hub_servername -file hub_cert256.arm
```

3. Restart the hub server and the spoke server.
4. Complete the following steps to define the spoke server to the hub server, and the hub server to the spoke server:
 - a. Issue the following commands on both the hub server and the spoke server:

```
SET SERVERPASSWORD server_password  
SET SERVERHLADDRESS ip_address  
SET SERVERLLADDRESS tcp_port
```

- b. On the hub server, issue the DEFINE SERVER command, according to the following example:

```
DEFINE SERVER spoke_servername HLA=spoke_address  
LLA=spoke_SSLTCPADMINPort SERVERPA=spoke_serverpassword
```

- c. On the spoke server, issue the DEFINE SERVER command, according to the following example:

```
DEFINE SERVER hub_servername HLA=hub_address  
LLA=hub_SSLTCPADMINPort SERVERPA=hub_serverpassword
```

Tip: By default, server communication is encrypted except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure. To encrypt all communication with the

specified server, even when the server is sending and receiving object data, specify the SSL=YES parameter on the DEFINE SERVER command.

5. Complete the following steps to configure the Operations Center to monitor the spoke server:
 - a. On the Operations Center menu bar, click Servers. The spoke server has a status of "Unmonitored." This status means that, although this server was defined to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke.
 - b. Click the spoke server to highlight the item, and click Monitor Spoke.

Related reference:

- [DEFINE SERVER \(Define a server for server-to-server communications\)](#)
- [QUERY OPTION \(Query server options\)](#)

Adding the second server as a spoke

After you configure both servers in your environment, add the second server as a spoke to the hub server.

Procedure

1. Open the Operations Center.
2. In the Operations Center menu bar, click Servers.
3. Complete one of the following steps:
 - o Click the server to highlight it, and in the table menu bar, click Monitor Spoke.
 - o If the server that you want to add is not shown in the table, click +Spoke.
4. Complete the steps in the spoke configuration wizard.

Enabling replication

To protect your data, enable node replication in addition to protecting your storage pools.

Procedure

To enable node replication for all of the clients that are registered to the source server, complete the following steps

1. Open the Operations Center.
2. On the Operations Center menu bar, hover over Storage and click Replication.
3. On the Replication page, click +Server Pair.
4. Complete the steps in the Add Server Pair wizard:
 - o Set the source server as the first server that you configured for the multisite disk solution. The target server is the second server.
 - o Set the node replication schedule to start 10 hours after the client backup window, based on the server maintenance activities that you scheduled in Defining schedules for server maintenance activities.
 - o The wizard sets up storage pool protection schedules for you, based on the amount of data that you are protecting and when client replication is scheduled.

What to do next

If you plan to set up mutual replication between the two sites, run the Add Server Pair wizard again and set the second server as the source and the first server as the target.

Completing the implementation

After the IBM Spectrum Protect™ solution is configured and running, test backup operations and set up monitoring to ensure that everything runs smoothly.

Procedure

1. Test backup operations to verify that your data is protected in the way that you expect.
 - a. On the Clients page of the Operations Center, select the clients that you want to back up, and click Back Up.

- b. On the Servers page of the Operations Center, select the server for which you want to back up the database. Click Back Up, and follow the instructions in the Back Up Database window.
 - c. Verify that the backup operations completed successfully with no warning or error messages.
Tip: Alternatively, you can use the backup-archive client GUI to back up client data and you can backup the server database by issuing BACKUP DB command from an administrative command-line.
2. Set up monitoring for your solution by following the instructions in Monitoring a multisite disk solution.

Monitoring a multisite disk solution

After you implement a multisite disk solution with IBM Spectrum Protect™, monitor the solution to ensure correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

About this task

The preferred way to monitor a solution is by using the Operations Center, which provides overall and detailed system status in a graphical user interface. In addition, you can configure the Operations Center to generate a daily email report that summarizes system status.

In some cases, you might want to use advanced monitoring tools to complete specific monitoring or troubleshooting tasks.

Tip: If you plan to diagnose issues with backup-archive clients on Linux or Windows operating systems, install IBM Spectrum Protect client management services on each computer where a backup-archive client is installed. In this way, you can ensure that the Diagnose button is available in the Operations Center for diagnosing issues with backup-archive clients. To install the client management service, follow the instructions in Installing the client management service.

Procedure

1. Complete daily monitoring tasks. For instructions, see Daily monitoring checklist.
2. Complete periodic monitoring tasks. For instructions, see Periodic monitoring checklist.
3. To verify that your IBM Spectrum Protect solution complies with licensing requirements, follow the instructions in Verifying license compliance.
4. To set up Operations Center to generate email status reports, see Tracking system status by using email reports

What to do next

Resolve any issues that you detect. To resolve an issue by changing the configuration of your solution, follow the instructions in Managing operations for a multisite disk solution. The following resources are also available:


- To resolve performance issues, see Performance.
- To resolve other types of issues, see Troubleshooting.


Daily monitoring checklist

To ensure that you are completing the daily monitoring tasks for your IBM Spectrum Protect™ solution, review the daily monitoring checklist.

Complete the daily monitoring tasks from the Operations Center Overview page. You can access the Overview page by opening the Operations Center and clicking Overviews.

The following figure shows the location for completing each task.








Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.


The following table lists the daily monitoring tasks and provides instructions for completing each task.

Table 1. Daily monitoring tasks



Task	Basic procedures	Advanced procedures and troubleshooting information
------	------------------	---




Task	Basic procedures	Advanced procedures and troubleshooting information
<p>Watch for security notifications, which can indicate a ransomware attack.</p>	<p>If a potential ransomware attack is detected in the IBM Spectrum Protect environment, a security notification message is displayed in the foreground of the Operations Center. For more information, click the message to open the Security Notifications page.</p>	<p>On the Security Notifications page, you can take the following actions:</p> <ul style="list-style-type: none"> • View notification details by client. Restriction: In Operations Center Version 8.1.5, notifications are available only for backup-archive clients. • Acknowledge a security notification by selecting it and clicking Acknowledge. When you acknowledge a security notification, a check mark is added to the Acknowledged column of the Security Notifications page for the selected client. The standard by which a notification is acknowledged is determined by your organization. A check mark might mean that you investigated the issue and determined that it is a false positive. Or it might mean that a problem exists and is being resolved. • Assign a security notification to an administrator by selecting the security notification and clicking Assign. To view the assignment, the administrator must sign in to the Operations Center and click Overviews > Security. If you are not certain whether the administrator regularly monitors the Security Notifications page, notify the administrator about the assignment. • If the notification is a false positive, you can select the security notification and click Reset. The security notification is deleted. Historical data that is used for baseline comparisons with the most recent backup operation is deleted. A new baseline is calculated going forward.
<p>1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.</p>	<p>To verify whether clients are at risk, in the Clients area, look for an At risk notification. To view details, click the Clients area. Attention: If the At risk percentage is much greater than usual, it might indicate a ransomware attack. A ransomware attack can cause backup operations to fail, thus placing clients at risk. For example, if the percentage of clients at risk is normally between 5% and 10%, but the percentage increases to 40% or 50%, investigate the cause. If you installed the client management service on a backup-archive client, you can view and analyze the client error and schedule logs by completing the following steps:</p> <ol style="list-style-type: none"> 1. In the Clients table, select the client and click Details. 2. To diagnose an issue, click Diagnosis. 	<p>For clients that do not have the client management service installed, access the client system to review the client error logs.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>2 Determine whether client-related or server-related errors require attention.</p>	<p>To determine the severity of any reported alerts, in the Alerts area, hover over the columns.</p>	<p>To view additional information about alerts, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Alerts area. 2. In the Alerts table, select an alert. 3. In the Activity Log pane, review the messages. The pane displays related messages that were issued before and after the selected alert occurred.
<p>3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.</p>	<ol style="list-style-type: none"> 1. To verify whether servers are at risk, in the Servers area, look for an Unavailable notification. 2. To view additional information, click the Servers area. 3. Select a server in the Servers table and click Details. 	<p>Tip: If you detect an issue that is related to server properties, update the server properties:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a server and click Details. 2. To update server properties, click Properties.
<p>4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.</p>	<ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Status column of the table, view the status of the server and resolve any issues: <ul style="list-style-type: none"> o Normal  Sufficient space is available for the server database, active log, and archive log. o Critical  Insufficient space is available for the server database, active log, or archive log. You must add space immediately, or the data protection services that are provided by the server will be interrupted. o Warning  The server database, active log, or archive log is running out of space. If this condition persists, you must add space. o Unavailable  Status cannot be obtained. Ensure that the server is running, and that there are no network issues. This status is also shown if the monitoring administrator ID is locked or otherwise unavailable on the server. This ID is named IBM-OC-hub_server_name. o Unmonitored  Unmonitored servers are defined to the hub server, but are not configured for management by the Operations Center. To configure an unmonitored server, select the server, and click Monitor Spoke. 	<p>You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>5 Verify server database backup operations.</p>	<p>To determine when a server was most recently backed up, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column. 	<p>To obtain more detailed information about backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a row and click Details. 2. In the DB Backup area, hover over the check marks to review information about backup operations. <p>If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click Back Up. <p>To determine whether the server database is configured for automatic backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the menu bar, hover over the settings icon  and click Command Builder. 2. Issue the QUERY DB command: <pre>query db f=d</pre> 3. In the output, review the Full Device Class Name field. If a device class is specified, the server is configured for automatic database backups.
<p>6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.</p>	<p>To search for information about processes that failed because of server issues, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click Servers > Maintenance. 2. To obtain the two-week history of a process, view the History column. 3. To obtain more information about a scheduled process, hover over the check box that is associated with the process. 	<p>For more information about monitoring processes and resolving issues, see the Operations Center online help.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>7 Verify that the amount of data that was recently sent to and from servers is within the expected range.</p>	<ul style="list-style-type: none"> To obtain an overview of activity in the last 24 hours, view the Activity area. To compare activity in the last 24 hours with activity in the previous 24 hours, review the figures in the Current and Previous areas. 	<ul style="list-style-type: none"> If more data was sent to the server than you expected, determine which clients are backing up more data and investigate the cause. It is possible that client-side data deduplication is not working correctly. Attention: If the amount of backed-up data is significantly larger than usual, it might indicate a ransomware attack. When ransomware encrypts data, the system perceives the data as being changed, and the changed data is backed up. Thus, backup volumes become larger. To determine which clients are affected, click the Applications, Virtual, or Systems tabs. If less data was sent to the server than you expected, investigate whether client backup operations are proceeding on schedule.

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>8 Verify that storage pools are available to back up client data.</p>	<ol style="list-style-type: none"> 1. If problems are indicated in the Storage & Data Availability area, click Pools to view the details: <ul style="list-style-type: none"> o If the Critical  status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. Attention: If the status is critical, investigate the cause: <ul style="list-style-type: none"> ■ If the data deduplication rate for a storage pool drops significantly, it might indicate a ransomware attack. During a ransomware attack, data is encrypted and cannot be deduplicated. To verify the data deduplication rate, in the Storage Pools table, review the value in the % Savings column. ■ If a storage pool unexpectedly becomes 100% utilized, it might indicate a ransomware attack. To verify the utilization, review the value in the Capacity Used column. Hover over the values to see the percentages of used and free space. o If the Warning  status is displayed, the storage pool is running out of space, or its access status is read-only. 2. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column. 	<p>To view the storage-pool capacity that was used over the past two weeks, select a row in the Storage Pools table and click Details.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>9 Verify that storage devices are available for backup operations.</p>	<p>In the Storage & Data Availability area, in the Volumes section, under the capacity bars, review the status that is reported next to Devices. If a Critical  or Warning  status is displayed for any device, investigate the issue. To view details, click Devices.</p>	<p>Disk devices might have a critical or warning status for the following reasons:</p> <ul style="list-style-type: none"> • For DISK device classes, volumes might be offline or have a read-only access status. The Disk Storage column of the Disk Devices table shows the state of volumes. • For FILE device classes that are not shared, directories might be offline. Also, insufficient free space might be available for allocating scratch volumes. The Disk Storage column of the Disk Devices table shows the state of directories. • For FILE device classes that are shared, drives might be unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. Other columns of the Disk Devices table show the state of the drives and paths.
<p>10 Monitor node replication processes.</p>	<ol style="list-style-type: none"> 1. To obtain the overall status of node replication processes, view the Replication area on the Operations Center Overview page. 2. To view information about each replicated server pair, click the Replication area. Attention: If you notice an unexpected increase in the number of replication failures, it might indicate a ransomware attack. Investigate the cause of the failures. 3. To view the amount of data that was replicated over the last two weeks and the speed of replication, select a server pair and click Details. 4. To view replication information for a client, on the Operations Center Overview page, click Clients. View the information in the Replication Workload column. Attention: If you see a drastic, unexpected increase in the replication workload, it might indicate a ransomware attack. Investigate the cause of the increased workload. 	<p>For advanced monitoring, view information about running and ended node replication processes by using commands:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. Issue the QUERY REPLICATION command. For instructions, see QUERY REPLICATION (Query node replication processes). If the replication operation was completed successfully, the <code>Total Files To Replicate</code> value matches the <code>Total Files Replicated</code> value. <p>To display messages that are related to a node replication process on a source or target replication server, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Servers. 2. Select the source or target replication server and click Details: <ul style="list-style-type: none"> ◦ To view active tasks, click Active Tasks, select the task, and verify that the Running status is displayed. For details, view the related activity logs. ◦ To view completed tasks, click Completed Tasks, select the task, and ensure that the Completed status is displayed. For details, view the related activity logs.

Periodic monitoring checklist

To help ensure that your solution operates correctly, complete the tasks in the periodic monitoring checklist. Schedule periodic tasks frequently enough so that you can detect potential issues before they become problematic.




Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.


Table 1. Periodic monitoring tasks

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor system performance.	<p>Determine the length of time that is required for client backup operations:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. Find the server that is associated with the client. 2. Click Servers. Select the server and click Details. 3. To view the duration of completed tasks in the last 24 hours, click Completed Tasks. 4. To view the duration of tasks that were completed more than 24 hours ago, use the QUERY ACTLOG command. Follow the instructions in . 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. <p>If you installed the client management service on a backup-archive client, you can diagnose performance issues for the backup-archive client by completing the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. 2. Select a backup-archive client and click Details. 3. To retrieve client logs, click Diagnosis. 	<p>For instructions about reducing the time that it takes for the client to back up data to the server, see Resolving common client performance problems.</p> <p>Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks.</p> <p>For information about identifying and resolving other performance issues, see Performance.</p>
Determine the disk savings that are provided by data deduplication.	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Pools. 2. Select a pool and click Quick Look. 3. In the Data Deduplication area, view the Space saved row. 	<p>For advanced monitoring, to obtain detailed statistics about the data-deduplication process for a specific directory-container storage pool or cloud-container storage pool, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. Obtain a statistical report by issuing the GENERATE DEDUPSTATS command. Follow the instructions in GENERATE DEDUPSTATS (Generate data deduplication statistics for a directory-container storage pool). 3. View the statistical report by issuing the QUERY DEDUPSTATS command. Follow the instructions in QUERY DEDUPSTATS (Query data deduplication statistics).

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Verify that current backup files for device configuration and volume history information are saved.</p>	<p>Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations.</p> <p>To locate the volume history and device configuration files, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To locate the volume history and device configuration files, issue the following commands: <ul style="list-style-type: none"> <code>query option volhistory</code> <code>query option devconfig</code> 3. In the output, review the Option Setting column to find the file locations. <p>If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.</p>	

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Determine whether sufficient space is available for the instance directory file system.</p>	<p>Verify that at least 20% of free space is available in the instance directory file system. Take the action that is appropriate for your operating system:</p> <ul style="list-style-type: none"> <p>AIX To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -g instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Linux To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -h instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Windows In the Windows Explorer program, right-click the file system and click Properties. View the capacity information.</p> <p>The preferred location of the instance directory depends on the operating system where the server is installed:</p> <ul style="list-style-type: none"> <p>AIX Linux /home/tsminst1/tsminst1</p> <p>Windows C:\tsminst1</p> <p>Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.</p>	
<p>Identify unexpected client activity.</p>	<p>To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps:</p> <ol style="list-style-type: none"> On the Operations Center Overview page, click the Clients area. To view activity over the past two weeks, double-click any client. To view the number of bytes sent to the client, click the Properties tab. In the Last Session area, view the Sent to client row. 	<p>When you double-click a client in the Clients table, the Activity over 2 Weeks area displays the amount of data that the client sent to the server each day.</p> <p>Periodically review the SQL activity summary table, which contains statistics about client sessions. To compare current activity with previous activity, use an SQL SELECT statement. If the level of activity is significantly different from previous activity, it might indicate a ransomware attack.</p> <p>Periodically review the activity log. Look for ANE messages that indicate how many files were backed up and inspected. Compare current data deduplication rates with previous rates. If an unusually high number of files were backed up, or the rate of data deduplication unexpectedly drops to 0, it might indicate a ransomware attack.</p>

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Monitor storage pool growth over time.</p>	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Pools area. 2. To view the capacity that was used over the last two weeks, select a pool and click Details. 	<p>Tips:</p> <ul style="list-style-type: none"> • To specify the time period that must elapse before all deduplicated extents are removed from a directory-container storage pool or cloud-container storage pool after they are no longer referenced by the inventory, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. 3. Specify the duration in the <code>Delay period for container reuse</code> field. • To determine data deduplication performance for directory-container and cloud-container storage pools, use the <code>GENERATE DEDUPSTATS</code> command. • To view data deduplication statistics for a storage pool, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. <p>Alternatively, use the <code>QUERY EXTENTUPDATES</code> command to display information about updates to data extents in directory-container or cloud-container storage pools. The command output can help you determine which data extents are no longer referenced and which ones are eligible to be deleted from the system. In the output, monitor the number of data extents that are eligible to be deleted from the system. This metric has a direct correlation to the amount of free space that will be available within the container storage pool.</p> <ul style="list-style-type: none"> • To display the amount of physical space that is occupied by a file space after the removal of the data deduplication savings, use the <code>select * from occupancy</code> command. The command output includes the <code>LOGICAL_MB</code> value. <code>LOGICAL_MB</code> is the amount of space that is used by the file space.
<p>Evaluate the timing of client schedules. Ensure that the start and end times of client schedules meet your business needs.</p>	<p>On the Operations Center Overview page, click Clients > Schedules.</p> <p>In the Schedules table, the Start column displays the configured start time for the scheduled operation. To see when the most recent operation was started, hover over the clock icon.</p>	<p>Tip: You can receive a warning message if a client operation runs longer than expected. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over Clients and click Schedules. 2. Select a schedule and click Details. 3. View the details of a schedule by clicking the blue arrow next to the row. 4. In the Run time alert field, specify the time when a warning message will be issued if the scheduled operation is not completed. 5. Click Save.

Task	Basic procedures	Advanced procedures and troubleshooting
Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks meet your business needs.	<p>On the Operations Center Overview page, click Servers > Maintenance.</p> <p>In the Maintenance table, review the information in the Last Run Time column. To see when the last maintenance task was started, hover over the clock icon.</p>	<p>Tip: If a maintenance task is running too long, change the start time or the maximum run time. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To change the start time or maximum run time for a task, issue the UPDATE SCHEDULE command. For instructions, see UPDATE SCHEDULE (Update a client schedule).

Related reference:

- [QUERY ACTLOG](#) (Query the activity log)
- [UPDATE STGPOOL](#) (Update a storage pool)
- [QUERY EXTENTUPDATES](#) (Query updated data extents)

Verifying license compliance

Verify that your IBM Spectrum Protect™ solution complies with the provisions of your licensing agreement. By verifying compliance regularly, you can track trends in data growth or processor value unit (PVU) usage. Use this information to plan for future license purchasing.

About this task

The method that you use to verify that your solution complies with license terms varies depending on the provisions of your IBM Spectrum Protect licensing agreement.

Front-end capacity licensing

The front-end model determines license requirements based on the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems.

Back-end capacity licensing

The back-end model determines license requirements based on the terabytes of data that are stored in primary storage pools and repositories.

Tips:

- To ensure the accuracy of front-end and back-end capacity estimates, install the most recent version of the client software on each client node.
- The front-end and back-end capacity information in the Operations Center is for planning and estimation purposes.

PVU licensing

The PVU model is based on the use of PVUs by server devices.

Important: The PVU calculations that are provided by IBM Spectrum Protect are considered estimates and are not legally binding. The PVU licensing information that is reported by IBM Spectrum Protect is not considered an acceptable substitute for the IBM® License Metric Tool.



For the most recent information about licensing models, see the information about product details and licenses at the IBM Spectrum Protect product family website. If you have questions or concerns about licensing requirements, contact your IBM Spectrum Protect software provider.

Procedure

To monitor license compliance, complete the steps that correspond to the provisions of your licensing agreement.

Tip: The Operations Center provides an email report that summarizes front-end and back-end capacity usage. Reports can be sent automatically to one or more recipients regularly. To configure and manage email reports, click Reports on the Operations Center menu bar.

Option	Description
--------	-------------

Option	Description
Front-end model	<p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>The front-end capacity estimate is displayed on the Front-end Usage page.</p> <p>b. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage.</p> <p>c. To estimate capacity for clients that did not report capacity usage, go to the following FTP site, which provides measuring tools and instructions:</p> <p><code>ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools</code></p> <p>To measure front-end capacity by script, complete the instructions in the most recently available licensing guide.</p> <p>d. Add the Operations Center estimate and any estimates that you obtained by using a script.</p> <p>e. Verify that the estimated capacity complies with your licensing agreement.</p>
Back-end model	<p>Restriction: If the source and target replication servers do not use the same policy settings, you cannot use the Operations Center to monitor back-end capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>b. Click the Back-end tab.</p> <p>c. Verify that the estimated amount of data complies with your licensing agreement.</p>
PVU model	<p>For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model.</p>

Tracking system status by using email reports

Set up the Operations Center to generate email reports that summarize system status. You can configure a mail server connection, change report settings, and optionally create custom reports.

Before you begin

Before you set up email reports, ensure that the following requirements are met:

- A Simple Mail Transfer Protocol (SMTP) host server is available to send and receive reports by email. The SMTP server must be configured as an open mail relay. You must also ensure that the IBM Spectrum Protect™ server that sends email messages has access to the SMTP server. If the Operations Center is installed on a separate computer, that computer does not require access to the SMTP server.
- To set up email reports, you must have system privilege for the server.
- To specify the recipients, you can enter one or more email addresses or administrator IDs. If you plan to enter an administrator ID, the ID must be registered on the hub server and must have an email address that is associated with it. To specify an email address for an administrator, use the EMAILADDRESS parameter of the UPDATE ADMIN command.

About this task

You can configure the Operations Center to send a general operations report, a license compliance report, and one or more custom reports. You create custom reports by selecting a template from a set of commonly used report templates or by entering SQL SELECT statements to query managed servers.

Procedure

To set up and manage email reports, complete the following steps:

1. On the Operations Center menu bar, click Reports.
2. If an email server connection is not yet configured, click Configure Mail Server and complete the fields. After you configure the mail server, the general operations report and license compliance report are enabled.

3. To change report settings, select a report, click Details, and update the form.
4. Optional: To add a custom report, click + Report, and complete the fields.
Tip: To immediately run and send a report, select the report and click Send.

Results

Enabled reports are sent according to the specified settings.

Related reference:

[UPDATE ADMIN](#) (Update an administrator)

Managing operations for a multisite disk solution

Use this information to manage operations for a multisite disk solution with IBM Spectrum Protect™ that includes a server and uses data deduplication for multiple locations.

- **Managing the Operations Center**
The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.
- **Protecting applications, virtual machines, and systems**
The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.
- **Managing data storage**
Manage your data for efficiency and add supported devices and media to the server to store client data.
- **Managing replication**
Use replication to recover data at a disaster recovery site and to maintain the same level of files on the source and target servers. You can manage replication at the node level. You can also protect data at the storage-pool level.
- **Securing the server**
Secure the IBM Spectrum Protect server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.
- **Stopping and starting the server**
Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.
- **Planning to upgrade the server**
When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.
- **Preparing for an outage or system update**
Prepare IBM Spectrum Protect to maintain your system in a consistent state during a planned power outage or system update.
- **Implementing a disaster recovery plan**
Implement a disaster recovery strategy to recover your applications if a disaster occurs and to ensure high server availability.
- **Recovering from data loss or system outages**
You can use IBM Spectrum Protect to recover data that was lost when a disaster or system outage occurred. You can recover directory-container storage pools, client data, and databases.

Managing the Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect™ environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

- **Adding and removing spoke servers**
In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.
- **Starting and stopping the web server**
The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.
- **Restarting the initial configuration wizard**
You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.

- Changing the hub server
You can use the Operations Center to remove the hub server of IBM Spectrum Protect, and configure another hub server.
- Restoring the configuration to the preconfiguration state
If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect servers are not defined as hub or spoke servers.

Adding and removing spoke servers

In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

About this task

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

- Adding a spoke server
After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.
- Removing a spoke server
You can remove a spoke server from the Operations Center.

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of the hub server.

Procedure

1. In the Operations Center menu bar, click Servers. The Servers page opens.

In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke server.

2. Complete one of the following steps:
 - Click the server to highlight it, and in the table menu bar, click Monitor Spoke.
 - If the server that you want to add is not shown in the table, and secure SSL/TLS communication is not required, click + Spoke in the table menu bar.
3. Provide the necessary information, and complete the steps in the spoke configuration wizard.
Tip: If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Removing a spoke server

You can remove a spoke server from the Operations Center.

About this task

You might need to remove a spoke server in the following situations, for example:

- You want to move the spoke server from one hub server to another hub server.
- You want to decommission the spoke server.

Procedure

To remove the spoke server from the group of servers that are managed by the hub server, complete the following steps:

1. From the IBM Spectrum Protect™ command line, issue the following command on the hub server:

```
QUERY MONITORSETTINGS
```

2. From the output of the command, copy the name that is in the Monitored Group field.
3. Issue the following command on the hub server, where *group_name* represents the name of the monitored group, and *member_name* represents the name of the spoke server:

```
DELETE GRPMEMBER group_name member_name
```

4. Optional: If you want to move the spoke server from one hub server to another hub server, do **not** complete this step. Otherwise, you can disable alerting and monitoring on the spoke server by issuing the following commands on the spoke server:

```
SET STATUSMONITOR OFF  
SET ALERTMONITOR OFF
```

5. Optional: If the spoke server definition is used for other purposes, such as enterprise configuration, command routing, storing virtual volumes, or library management, do **not** complete this step. Otherwise, you can delete the spoke server definition on the hub server by issuing the following command on the hub server:

```
DELETE SERVER spoke_server_name
```

Tip: If a server definition is deleted immediately after the server is removed from the monitored group, status information for the server can remain in the Operations Center indefinitely.

To avoid this issue, wait until the status collection interval passes before you delete the server definition. The status collection interval is shown on the Settings page of the Operations Center.

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.

Procedure

1. Stop the web server.
 - o **AIX** From the */installation_dir/ui/utls* directory, where *installation_dir* represents the directory where the Operations Center is installed, issue the following command:

```
./stopserver.sh
```
 - o **Linux** Issue the following command:

```
service opscenter.rc stop
```
 - o **Windows** From the Services window, stop the IBM Spectrum Protect™ Operations Center service.
2. Start the web server.
 - o **AIX** From the */installation_dir/ui/utls* directory, where *installation_dir* represents the directory where the Operations Center is installed, issue the following command:

```
./startserver.sh
```
 - o **Linux** Issue the following commands:

Start the server:

```
service opscenter.rc start
```

Restart the server:

```
service opscenter.rc restart
```

Determine whether the server is running:

```
service opscenter.rc status
```
 - o **Windows** From the Services window, start the IBM Spectrum Protect Operations Center service.

Restarting the initial configuration wizard

You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.

Before you begin

To change the following settings, use the Settings page in the Operations Center rather than restarting the initial configuration wizard:

- The frequency at which status data is refreshed
- The duration that alerts remain active, inactive, or closed
- The conditions that indicate that clients are at risk

The Operations Center help includes more information about how to change these settings.

About this task

To restart the initial configuration wizard, you must delete a properties file that includes information about the hub server connection. However, any alerting, monitoring, at-risk, or multiserver settings that were configured for the hub server are not deleted. These settings are used as the default settings in the configuration wizard when the wizard restarts.

Procedure

1. Stop the Operations Center web server.
 2. On the computer where the Operations Center is installed, go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:
 - o **AIX** | **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer`
 - o **Windows** `installation_dir\ui\Liberty\usr\servers\guiServer`
- For example:
- o **AIX** | **Linux** `/opt/tivoli/tsm/ui/Liberty/usr/servers/guiServer`
 - o **Windows** `c:\Program Files\Tivoli\TSM\ui\Liberty\usr\servers\guiServer`
3. In the `guiServer` directory, delete the `serverConnection.properties` file.
 4. Start the Operations Center web server.
 5. Open the Operations Center.
 6. Use the configuration wizard to reconfigure the Operations Center. Specify a new password for the monitoring administrator ID.
 7. On any spoke servers that were previously connected to the hub server, update the password for the monitoring administrator ID by issuing the following command from the IBM Spectrum Protect™ command-line interface:

```
UPDATE ADMIN IBM-OC-hub_server_name new_password
```

Restriction: Do not change any other settings for this administrator ID. After you specify the initial password, this password is managed automatically by the Operations Center.

Changing the hub server

You can use the Operations Center to remove the hub server of IBM Spectrum Protect™, and configure another hub server.

Procedure

1. Restart the initial configuration wizard of the Operations Center. As part of this procedure, you delete the existing hub server connection.
2. Use the wizard to configure the Operations Center to connect to the new hub server.

Related tasks:

Restarting the initial configuration wizard

Restoring the configuration to the preconfiguration state

If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect™ servers are not defined as hub or spoke servers.

Procedure

To restore the configuration, complete the following steps:

1. Stop the Operations Center web server.
2. Unconfigure the hub server by completing the following steps:
 - a. On the hub server, issue the following commands:

```
SET MONITORINGADMIN ""
SET MONITOREDSEVERGROUP ""
SET STATUSMONITOR OFF
SET ALERTMONITOR OFF
REMOVE ADMIN IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. Reset the password for the hub server by issuing the following command on the hub server:

```
SET SERVERPASSWORD ""
```

Attention: Do not complete this step if the hub server is configured with other servers for other purposes, such as library sharing, exporting and importing of data, or node replication.

3. Unconfigure any spoke servers by completing the following steps:
 - a. On the hub server, to determine whether any spoke servers remain as members of the server group, issue the following command:

```
QUERY SERVERGROUP IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the name of the monitored server group that was automatically created when you configured the first spoke server. This server group name is also the same as the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. On the hub server, to delete spoke servers from the server group, issue the following command for each spoke server:

```
DELETE GRPMEMBER IBM-OC-hub_server_name spoke_server_name
```

- c. After all spoke servers are deleted from the server group, issue the following commands on the hub server:

```
DELETE SERVERGROUP IBM-OC-hub_server_name
SET MONITOREDSEVERGROUP ""
```

- d. On each spoke server, issue the following commands:

```
REMOVE ADMIN IBM-OC-hub_server_name
SETOPT PUSHSTATUS NO
SET ALERTMONITOR OFF
SET STATUSMONITOR OFF
```

- e. On each spoke server, delete the definition of the hub server by issuing the following command:

```
DELETE SERVER hub_server_name
```

Attention: Do not complete this step if the definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

- f. On the hub server, delete the definition of each spoke server by issuing the following command:

```
DELETE SERVER spoke_server_name
```

Attention: Do not complete this step if the server definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

4. Restore the default settings on each server by issuing the following commands:

```
SET STATUSREFRESHINTERVAL 5
SET ALERTUPDATEINTERVAL 10
SET ALERTACTIVEDURATION 480
SET ALERTINACTIVEDURATION 480
SET ALERTCLOSEDDURATION 60
```

```
SET STATUSATRISKINTERVAL TYPE=AP INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=VM INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=SY INTERVAL=24
SET STATUSSKIPASFAILURE YES TYPE=ALL
```

5. Restart the initial configuration wizard of the Operations Center.

Related tasks:

Restarting the initial configuration wizard
Starting and stopping the web server

Protecting applications, virtual machines, and systems

The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.

- **Adding clients**
After you implement a data protection solution with IBM Spectrum Protect, you can expand the solution by adding clients.
- **Managing client operations**
You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.
- **Managing client upgrades**
When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.
- **Decommissioning a client node**
If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect server, but the workstation is no longer used, you can decommission the workstation.
- **Deactivating data to free storage space**
In some cases, you can deactivate data that is stored on the IBM Spectrum Protect server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

Adding clients

After you implement a data protection solution with IBM Spectrum Protect™, you can expand the solution by adding clients.

About this task

The procedure describes basic steps for adding a client. For more specific instructions about configuring clients, see the documentation for the product that you install on the client node. You can have the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

Procedure

To add a client, complete the following steps:

1. Select the software to install on the client node and plan the installation. Follow the instructions in [Selecting the client software and planning the installation](#).
2. Specify how to back up and archive client data. Follow the instructions in [Specifying rules for backing up and archiving client data](#).
3. Specify when to back up and archive client data. Follow the instructions in [Scheduling backup and archive operations](#).
4. To allow the client to connect to the server, register the client. Follow the instructions in [Registering clients](#).
5. To start protecting a client node, install and configure the selected software on the client node. Follow the instructions in [Installing and configuring clients](#).

Selecting the client software and planning the installation

Different types of data require different types of protection. Identify the type of data that you must protect and select the appropriate software.

About this task

The preferred practice is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you install a product for which the client acceptor does not run schedules, you must follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

Procedure

Based on your goal, select the products to install and review the installation instructions.

Tip: If you install the client software now, you must also complete the client configuration tasks that are described in [Installing and configuring clients](#) before you can use the client.

Goal	Product and description	Installation instructions
Protect a file server or workstation	The backup-archive client backs up and archives files and directories from file servers and workstations to storage. You can also restore and retrieve backup versions and archived copies of files.	<ul style="list-style-type: none"> • Backup-archive client requirements • Install UNIX and Linux backup-archive clients • Installing the Windows client for the first time
Protect applications with snapshot backup and restore capabilities	IBM Spectrum Protect Snapshot protects data with integrated, application-aware snapshot backup and restore capabilities. You can protect data that is stored by IBM DB2® database software and SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications.	<ul style="list-style-type: none"> • Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux • Installing and upgrading IBM Spectrum Protect Snapshot for VMware • Installing and upgrading IBM Spectrum Protect Snapshot for Windows
Protect an email application on an IBM Domino® server	IBM Spectrum Protect for Mail: Data Protection for IBM® Domino automates data protection so that backups are completed without shutting down IBM Domino servers.	<ul style="list-style-type: none"> • Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) • Installation of Data Protection for IBM Domino on a Windows system (V7.1.0)
Protect an email application on a Microsoft Exchange server	IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server automates data protection so that backups are completed without shutting down Microsoft Exchange servers.	Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
Protect an IBM DB2 database	The application programming interface (API) of the backup-archive client can be used to back up DB2 data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)

Goal	Product and description	Installation instructions
Protect an IBM Informix® database	The API of the backup-archive client can be used to back up Informix data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)
Protect a Microsoft SQL database	IBM Spectrum Protect for Databases: Data Protection for Microsoft SQL Server protects Microsoft SQL data.	Installing Data Protection for SQL Server on Windows Server Core
Protect an Oracle database	IBM Spectrum Protect for Databases: Data Protection for Oracle protects Oracle data.	Data Protection for Oracle installation
Protect an SAP environment	IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP provides protection that is customized for SAP environments. The product is designed to improve the availability of SAP database servers and reduce administration workload.	<ul style="list-style-type: none"> Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2 Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
Protect a virtual machine	<p>IBM Spectrum Protect for Virtual Environments provides protection that is tailored for Microsoft Hyper-V and VMware virtual environments. You can use IBM Spectrum Protect for Virtual Environments to create incremental forever backups that are stored on a centralized server, create backup policies, and restore virtual machines or individual files.</p> <p>Alternatively, use the backup-archive client to back up and restore a full VMware or Microsoft Hyper-V virtual machine. You can also back up and restore files or directories from a VMware virtual machine.</p>	<ul style="list-style-type: none"> Installing Data Protection for Microsoft Hyper-V Installing and upgrading Data Protection for VMware Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)

Tip: To use the client for space management, you can install IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows.

Specifying rules for backing up and archiving client data

Before you add a client, ensure that appropriate rules are specified for backup and archive operations for the client data. During the client registration process, you assign the client node to a policy domain, which has the rules that control how and when client data is stored.

Before you begin

Determine how to proceed:

- If you are familiar with the policies that are configured for your solution and you know that they do not require changes, continue with Scheduling backup and archive operations.
- If you are not familiar with the policies, follow the steps in this procedure.

About this task

Policies affect how much data is stored over time, and how long data is retained and available for clients to restore. To meet objectives for data protection, you can update the default policy and create your own policies. A policy includes the following rules:

- How and when files are backed up and archived to server storage.
- The number of copies of a file and the length of time copies are kept in server storage.

During the client registration process, you assign a client to a *policy domain*. The policy for a specific client is determined by the rules in the policy domain to which the client is assigned. In the policy domain, the rules that are in effect are in the active *policy*

set.

When a client backs up or archives a file, the file is bound to a management class in the active policy set of the policy domain. A *management class* is the key set of rules for managing client data. The backup and archive operations on the client use the settings in the default management class of the policy domain unless you further customize policy. A policy can be customized by defining more management classes and assigning their use through client options.

Client options can be specified in a local, editable file on the client system and in a client option set on the server. The options in the client option set on the server can override or add to the options in the local client option file.

Procedure

1. Review the policies that are configured for your solution by following the instructions in Viewing policies.
2. If you need to make minor changes to meet data retention requirements, follow the instructions in Editing policies.
3. Optional: If you need to create policy domains or make extensive changes to policies to meet data retention requirements, see Customizing policies.

Viewing policies

View policies to determine whether they must be edited to meet your requirements.

Procedure

1. To view the active policy set for a policy domain, complete the following steps:
 - a. On the Services page of the Operations Center, select a policy domain and click Details.
 - b. On the Summary page for the policy domain, click the Policy Sets tab.

Tip: To help ensure that you can recover data after a ransomware attack, apply the following guidelines:

 - Ensure that the value in the Backups column is a minimum of 2. The preferred value is 3, 4, or more.
 - Ensure that the value in the Keep Extra Backups column is a minimum of 14 days. The preferred value is 30 or more days.
 - Ensure that the value in the Keep Archives column is a minimum of 30 days.

If IBM Spectrum Protect™ for Space Management software is installed on the client, ensure that data is backed up before you migrate it. On the DEFINE MGMTCLASS or UPDATE MGMTCLASS command, specify MIGQUIRESBKUP=YES. Then, follow the guidelines in the tip.
2. To view inactive policy sets for a policy domain, complete the following steps:
 - a. On the Policy Sets page, click the Configure toggle. You can now view and edit the policy sets that are inactive.
 - b. Scroll through the inactive policy sets by using the forward and back arrows. When you view an inactive policy set, the settings that differentiate the inactive policy set from the active policy set are highlighted.
 - c. Click the Configure toggle. The policy sets are no longer editable.

Editing policies

To change the rules that apply to a policy domain, edit the active policy set for the policy domain. You can also activate a different policy set for a domain.

Before you begin

Changes to policy can affect data retention. Ensure that you continue to back up data that is essential to your organization so that you can restore that data if a disaster occurs. Also, ensure that your system has sufficient storage space for planned backup operations.

About this task

You edit a policy set by changing one or more management classes within the policy set. If you edit the active policy set, the changes are not available to clients unless you reactivate the policy set. To make the edited policy set available to clients, activate the policy set.

Although you can define multiple policy sets for a policy domain, only one policy set can be active. When you activate a different policy set, it replaces the currently active policy set.

To learn about preferred practices for defining policies, see Customizing policies.

Procedure

1. On the Services page of the Operations Center, select a policy domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab.

The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.

3. Click the Configure toggle. The policy set is editable.
4. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.
5. Edit the policy set by completing any of the following actions:

Option	Description
Add a management class	<ol style="list-style-type: none">a. In the Policy Sets table, click +Management Class.b. To specify the rules for backing up and archiving data, complete the fields in the Add Management Class window.c. To make the management class the default management class, select the Make default check box.d. Click Add.
Delete a management class	In the Management Class column, click -. Tip: To delete the default management class, you must first assign a different management class as the default.
Make a management class the default management class	In the Default column for the management class, click the radio button. Tip: The default management class manages client files when another management class is not assigned to, or appropriate for managing, a file. To ensure that clients can always back up and archive files, choose a default management class that contains rules for both backing up and archiving files.
Modify a management class	To change the properties of a management class, update the fields in the table.

6. Click Save.
Attention: When you activate a new policy set, data might be lost. Data that is protected under one policy set might not be protected under another policy set. Therefore, before you activate a policy set, ensure that the differences between the previous policy set and the new policy set do not cause data to be lost.
7. Click Activate. A summary of the differences between the active policy set and the new policy set is displayed. Ensure that the changes in the new policy set are consistent with your data retention requirements by completing the following steps:
 - a. Review the differences between corresponding management classes in the two policy sets, and consider the consequences for client files. Client files that are bound to management classes in the active policy set will be bound to the management classes with the same names in the new policy set.
 - b. Identify management classes in the active policy set that do not have counterparts in the new policy set, and consider the consequences for client files. Client files that are bound to these management classes will be managed by the default management class in the new policy set.
 - c. If the changes to be implemented by the policy set are acceptable, select the I understand that these updates can cause data loss check box and click Activate.

Scheduling backup and archive operations

Before you register a new client with the server, ensure that a schedule is available to specify when backup and archive operations take place. During the registration process, you assign a schedule to the client.

Before you begin

Determine how to proceed:

- If you are familiar with the schedules that are configured for the solution and you know that they do not require modification, continue with Registering clients.
- If you are not familiar with the schedules or the schedules require modification, follow the steps in this procedure.

About this task


Typically, backup operations for all clients must be completed daily. Schedule client and server workloads to achieve the best performance for your storage environment. To avoid the overlap of client and server operations, consider scheduling client backup

and archive operations so that they run at night. If client and server operations overlap or are not given enough time and resources to be processed, you might experience decreased system performance, failed operations, and other issues.

Procedure

1. Review available schedules by hovering over Clients on the Operations Center menu bar. Click Schedules.
2. Optional: Modify or create a schedule by completing the following steps:

Option	Description
Modify a schedule	<ol style="list-style-type: none">a. In the Schedules view, select the schedule and click Details.b. On the Schedule Details page, view details by clicking the blue arrows at the beginning of the rows.c. Modify the settings in the schedule, and click Save.
Create a schedule	In the Schedules view, click +Schedule and complete the steps to create a schedule.

3. Optional: To configure schedule settings that are not visible in the Operations Center, use a server command. For example, you might want to schedule a client operation that backs up a specific directory and assigns it to a management class other than the default.
 - a. On the Overview page of the Operations Center, hover over the settings icon  and click Command Builder.
 - b. Issue the DEFINE SCHEDULE command to create a schedule or the UPDATE SCHEDULE command to modify a schedule. For more information about the commands, see DEFINE SCHEDULE (Define a schedule for an administrative command) or UPDATE SCHEDULE (Update a client schedule).

Related tasks:

[Tuning the schedule for daily operations](#)

Registering clients

Register a client to ensure that the client can connect to the server, and the server can protect client data.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see technote 7048963.

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see Authenticating users by using an Active Directory database.

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the REGISTER NODE command and specify the USERID parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see Register a node.

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:
 - a. On the Operations Center menu bar, click Clients.
 - b. In the Clients table, click +Client.
 - c. Complete the steps in the Add Client wizard:
 - i. Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the Enable check box.
 - ii. In the Configuration window, copy the TCPSEVERADDRESS, TCPPOINT, NODENAME, and DEDUPLICATION option values.
Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
 - iii. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.

- iv. Set how risks are displayed for the client by specifying the at-risk setting.
- v. Click Add Client.

Related reference:

- [Tcpserveraddress option](#)
- [Tcpsport option](#)
- [Nodename option](#)
- [Deduplication option](#)

Installing and configuring clients

To start protecting a client node, you must install and configure the selected software.

Procedure

If you already installed the software, start at step 2.

1. Take one of the following actions:
 - o To install software on an application or client node, follow the instructions.

Software	Link to instructions
IBM Spectrum Protect™ backup-archive client	<ul style="list-style-type: none"> ▪ Install UNIX and Linux backup-archive clients ▪ Installing the Windows client for the first time <p>Tip: You can also update existing clients by using the Operations Center. For instructions, see Scheduling client updates.</p>
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> ▪ Data Protection for Oracle installation ▪ Installing Data Protection for SQL Server on Windows Server Core
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> ▪ Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) ▪ Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) ▪ Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> ▪ Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux ▪ Installing and upgrading IBM Spectrum Protect Snapshot for VMware ▪ Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle

- o To install software on a virtual machine client node, follow the instructions for the selected backup type.

Backup type	Link to instructions
If you plan to create full VMware backups of virtual machines, install and configure the IBM Spectrum Protect backup-archive client.	<ul style="list-style-type: none"> ▪ Install UNIX and Linux backup-archive clients ▪ Installing the Windows client for the first time
If you plan to create incremental forever full backups of virtual machines, install and configure IBM Spectrum Protect for Virtual Environments and the backup-archive client on the same client node or on different client nodes.	<ul style="list-style-type: none"> ▪ IBM Spectrum Protect for Virtual Environments online product documentation <p>Tip: You can obtain the software for IBM Spectrum Protect for Virtual Environments and the backup-archive client in the IBM Spectrum Protect for Virtual Environments installation package.</p>

2. To allow the client to connect to the server, add or update the values for the TCPSERVERADDRESS, TCPSPORT, and NODENAME options in the client options file. Use the values that you recorded when you registered the client (Registering clients).

- For clients that are installed on an AIX®, Linux, or Mac OS X operating system, add the values to the client system-options file, `dsm.sys`.
- For clients that are installed on a Windows operating system, add the values to the `dsm.opt` file.

By default, the options files are in the installation directory.

3. If you installed a backup-archive client on a Linux or Windows operating system, install the client management service on the client. Follow the instructions in [Installing the client management service](#).
4. Configure the client to run scheduled operations. Follow the instructions in [Configuring the client to run scheduled operations](#).
5. Optional: Configure communications through a firewall. Follow the instructions in [Configuring client/server communications through a firewall](#).
6. Run a test backup to verify that data is protected as you planned. For example, for a backup-archive client, complete the following steps:
 - a. On the **Clients** page of the Operations Center, select the client that you want to back up, and click **Back Up**.
 - b. Verify that the backup completes successfully and that there are no warning or error messages.
7. Monitor the results of the scheduled operations for the client in the Operations Center.

What to do next

To change what is getting backed up from the client, follow the instructions in [Modifying the scope of a client backup](#).

Configuring the client to run scheduled operations

You must configure and start a client scheduler on the client node. The client scheduler enables communication between the client and server so that scheduled operations can occur. For example, scheduled operations typically include backing up files from a client.

About this task

The preferred method is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations. The client acceptor manages the client scheduler so that the scheduler runs only when required:

- When it is time to query the server about the next scheduled operation
- When it is time to start the next scheduled operation

By using the client acceptor, you can reduce the number of background processes on the client and help to avoid memory retention problems.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you installed a product for which the client acceptor does not run schedules, follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

If your business uses a third-party scheduling tool as standard practice, you can use that scheduling tool as an alternative to the client acceptor. Typically, third-party scheduling tools start client programs directly by using operating system commands. To configure a third-party scheduling tool, see the product documentation.

Procedure

To configure and start the client scheduler by using the client acceptor, follow the instructions for the operating system that is installed on the client node:

AIX® and Oracle Solaris

- a. From the backup-archive client GUI, click **Edit > Client Preferences**.
- b. Click the **Web Client** tab.
- c. In the **Managed Services Options** field, click **Schedule**. If you also want the client acceptor to manage the web client, click the **Both** option.
- d. To ensure that the scheduler can start unattended, in the `dsm.sys` file, set the `passwordaccess` option to `generate`.
- e. To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

f. Start the client acceptor by issuing the following command on the command line:

```
/usr/bin/dsmcad
```

g. To enable the client acceptor to start automatically after a system restart, add the following entry to the system startup file (typically, /etc/inittab):

```
tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # Client Acceptor Daemon
```

Linux

- From the backup-archive client GUI, click Edit > Client Preferences.
- Click the Web Client tab.
- In the Managed Services Options field, click Schedule. If you also want the client acceptor to manage the web client, click the Both option.
- To ensure that the scheduler can start unattended, in the dsm.sys file, set the passwordaccess option to generate.
- To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

f. Start the client acceptor by logging in with the root user ID and issuing the following command:

```
service dsmcad start
```

g. To enable the client acceptor to start automatically after a system restart, add the service by issuing the following command at a shell prompt:

```
# chkconfig --add dsmcad
```

MAC OS X

- In the backup-archive client GUI, click Edit > Client Preferences.
- To ensure that the scheduler can start unattended, click Authorization, select Password Generate, and click Apply.
- To specify how services are managed, click Web Client, select Schedule, click Apply, and click OK.
- To ensure that the generated password is saved, restart the backup-archive client.
- Use the IBM Spectrum Protect Tools for Administrators application to start the client acceptor.

Windows

- In the backup-archive client GUI, click Utilities > Setup Wizard > Help me configure the Client Scheduler. Click Next.
- Read the information on the Scheduler Wizard page and click Next.
- On the Scheduler Task page, select Install a new or additional scheduler and click Next.
- On the Scheduler Name and Location page, specify a name for the client scheduler that you are adding. Then, select Use the Client Acceptor daemon (CAD) to manage the scheduler and click Next.
- Enter the name that you want to assign to this client acceptor. The default name is Client Acceptor. Click Next.
- Complete the configuration by stepping through the wizard.
- Update the client options file, dsm.opt, and set the passwordaccess option to generate.
- To store the client node password, issue the following command at the command prompt:

```
dsmc query sess
```

Enter the client node password when prompted.

- Start the client acceptor service from the Services Control page. For example, if you used the default name, start the Client Acceptor service. Do not start the scheduler service that you specified on the Scheduler Name and Location page. The scheduler service is started and stopped automatically by the client acceptor service as needed.

Configuring client/server communications through a firewall

If a client must communicate with a server through a firewall, you must enable client/server communications through the firewall.

Before you begin

If you used the Add Client wizard to register a client, find the option values in the client options file that you obtained during that process. You can use the values to specify ports.

About this task

Attention: Do not configure a firewall in a way that might cause termination of sessions that are in use by a server or storage agent. Termination of a valid session can cause unpredictable results. Processes and sessions might appear to stop due to input/output errors. To help exclude sessions from timeout restrictions, configure known ports for IBM Spectrum Protect™ components. Ensure that the KEEPALIVE server option remains set to the default value of YES. In this way, you can help to ensure that client/server communication is uninterrupted. For instructions about setting the KEEPALIVE server option, see KEEPALIVE.

Procedure

Open the following ports to allow access through the firewall:

TCP/IP port for the backup-archive client, command-line administrative client, and the client scheduler

Specify the port by using the `tcpport` option in the client options file. The `tcpport` option in the client options file must match the `TCPPOINT` option in the server options file. The default value is 1500. If you decide to use a value other than the default, specify a number in the range 1024 - 32767.

HTTP port to enable communication between the web client and remote workstations

Specify the port for the remote workstation by setting the `httpport` option in the client options file of the remote workstation. The default value is 1581.

TCP/IP ports for the remote workstation

The default value of 0 (zero) causes two free port numbers to be randomly assigned to the remote workstation. If you do not want the port numbers to be randomly assigned, specify values by setting the `webports` option in the client options file of the remote workstation.

TCP/IP port for administrative sessions

Specify the port on which the server waits for requests for administrative client sessions. The value of the client `tcpadminport` option must match the value of the `TCPADMINPORT` server option. In this way, you can secure administrative sessions within a private network.

Managing client operations

You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.

About this task

In some cases, you can resolve client errors by stopping and starting the client acceptor. If client nodes or administrator IDs are locked, you can resolve the issue by unlocking the client node or administrator ID, and then resetting the password.

For detailed instructions about identifying and resolving client errors, see [Resolving client problems](#).

- Evaluating errors in client error logs
You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.
- Stopping and restarting the client acceptor
If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.
- Resetting passwords
If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.
- Modifying the scope of a client backup
When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

Evaluating errors in client error logs

You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

Before you begin

To resolve errors in a backup-archive client on a Linux or Windows operating system, ensure that the client management service is installed and started. For installation instructions, see [Installing the client management service](#). For instructions about verifying the installation, see [Verifying that the client management service is installed correctly](#).

Procedure

To diagnose and resolve client errors, take one of the following actions:

- If the client management service is installed on the client node, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click Details.
 3. On the client Summary page, click the Diagnosis tab.
 4. Review the retrieved log messages.

Tips:

- To show or hide the Client Logs pane, double-click the Client Logs bar.
- To resize the Client Logs pane, click and drag the Client Logs bar.

If suggestions are displayed on the Diagnosis page, select a suggestion. In the Client Logs pane, client log messages to which the suggestion relates are highlighted.

5. Use the suggestions to resolve the problems that are indicated by the error messages.
Tip: Suggestions are provided for only a subset of client messages.
- If the client management service is not installed on the client node, review the error logs for the installed client.

Stopping and restarting the client acceptor

If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.

About this task

In some cases, you can resolve client scheduling problems by stopping and restarting the client acceptor. The client acceptor must be running to ensure that scheduled operations can occur on the client. For example, if you change the IP address or domain name of the server, you must restart the client acceptor.

Procedure

Follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

- To stop the client acceptor, complete the following steps:
 - a. Determine the process ID for the client acceptor by issuing the following command on the command line:

```
ps -ef | grep dsmcad
```

Review the output. In the following sample output, 6764 is the process ID for the client acceptor:

```
root 6764 1 0 16:26:35 ? 0:00 /usr/bin/dsmcad
```

- b. Issue the following command on the command line:

```
kill -9 PID
```

where *PID* specifies the process ID for the client acceptor.

- To start the client acceptor, issue the following command on the command line:

```
/usr/bin/dsmcad
```

Linux

- To stop the client acceptor (and not restart it), issue the following command:

```
# service dsmcad stop
```

- To stop and restart the client acceptor, issue the following command:

```
# service dsmcad restart
```

MAC OS X

Click Applications > Utilities > Terminal.

- To stop the client acceptor, issue the following command:

```
/bin/launchctl unload -w com.ibm.tivoli.dsmcad
```

- To start the client acceptor, issue the following command:

```
/bin/launchctl load -w com.ibm.tivoli.dsmcad
```

Windows

- To stop the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Stop and OK.
- To restart the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Start and OK.

Related reference:

[Resolving client scheduling problems](#)

Resetting passwords

If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

Procedure

To resolve password issues, take one of the following actions:

- If a backup-archive client is installed on a client node, and the password is lost or forgotten, complete the following steps:
 1. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwnreset=yes
```

where *node_name* specifies the client node and *new_password* specifies the password that you assign.

2. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to `generate` in the client options file.

- If an administrator is locked out because of password issues, complete the following steps:
 1. To provide the administrator with access to the server, issue the UNLOCK ADMIN command. For instructions, see UNLOCK ADMIN (Unlock an administrator).
 2. Set a new password by using the UPDATE ADMIN command:

```
update admin admin_name new_password forcepwnreset=yes
```

where *admin_name* specifies the name of the administrator and *new_password* specifies the password that you assign.

- If a client node is locked, complete the following steps:
 1. Determine why the client node is locked and whether it must be unlocked. For example, if the client node is decommissioned, the client node is being removed from the production environment. You cannot reverse the

decommission operation, and the client node remains locked. A client node also might be locked if the client data is the subject of a legal investigation.

2. If you must unlock a client node, use the UNLOCK NODE command. For instructions, see UNLOCK NODE (Unlock a client node).
3. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwreset=yes
```

where *node_name* specifies the name of the node and *new_password* specifies the password that you assign.

4. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to `generate` in the client options file.

Modifying the scope of a client backup

When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

About this task

When you exclude unnecessary objects from backup operations, you get better control of the amount of storage space that is required for backup operations, and the cost of storage. Depending on your licensing package, you also might be able to limit licensing costs.

Procedure

How you modify the scope of backup operations depends on the product that is installed on the client node:

- For a backup-archive client, you can create an include-exclude list to include or exclude a file, groups of files, or directories from backup operations. To create an include-exclude list, follow the instructions in [Creating an include-exclude list](#).

To ensure consistent use of an include-exclude list for all clients of one type, you can create a client option set on the server that contains the required options. Then, you assign the client option set to each of the clients of the same type. For details, see [Controlling client operations through client option sets](#).

- For a backup-archive client, you can specify the objects to include in an incremental backup operation by using the domain option. Follow the instructions in [Domain option](#).
- For other products, to define which objects are included in and excluded from backup operations, follow the instructions in the product documentation.

Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Before you begin

1. Review the client/server compatibility requirements in technote 1053218. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted.
2. Verify system requirements for the client in IBM Spectrum Protect™ Supported Operating Systems.
3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with servers that are configured as library managers. See technote 1302789.

If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, complete the instructions that are listed in the following table.

Software	Link to instructions
----------	----------------------

Software	Link to instructions
IBM Spectrum Protect backup-archive client	<ul style="list-style-type: none"> Scheduling client updates
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux Installing and upgrading IBM Spectrum Protect Snapshot for VMware Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> Upgrading Data Protection for SQL Server Data Protection for Oracle installation Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Virtual Environments	<ul style="list-style-type: none"> Installing and upgrading Data Protection for VMware Installing Data Protection for Microsoft Hyper-V

Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect™ server, but the workstation is no longer used, you can decommission the workstation.

About this task

When you start the decommission process, the server locks the client node to prevent it from accessing the server. Files that belong to the client node are gradually deleted, and then the client node is deleted. You can decommission the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

The simplest method for decommissioning a client node is to use the Operations Center. The decommission process runs in the background. If the client is configured to replicate client data, the Operations Center automatically removes the client from replication on the source and target replication servers before it decommissions the client.

Tip: Alternatively, you can decommission a client node by issuing the DECOMMISSION NODE or DECOMMISSION VM command. You might want to use this method in the following cases:

- To schedule the decommission process for the future or to run a series of commands by using a script, specify the decommission process to run in the background.
- To monitor the decommission process for debugging purposes, specify the decommission process to run in the foreground. If you run the process in the foreground, you must wait for the process to be completed before you continue with other tasks.

Procedure

Take one of the following actions:

- To decommission a client in the background by using the Operations Center, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click More > Decommission.
- To decommission a client node by using an administrative command, complete the following steps:
 1. Determine whether the client node is configured for node replication by issuing the QUERY NODE command. For example, if the client node is named AUSTIN, run the following command:

```
query node austin format=detailed
```

Review the Replication State output field.

2. If the client node is configured for replication, remove the client node from replication by issuing the REMOVE REPLNODE command. For example, if the client node is named AUSTIN, issue the following command:

```
remove replnode austin
```

3. Take one of the following actions:

- To decommission an application or system client node in the background, issue the DECOMMISSION NODE command. For example, if the client node is named AUSTIN, issue the following command:

```
decommission node austin
```

- To decommission an application or system client node in the foreground, issue the DECOMMISSION NODE command and specify the `wait=yes` parameter. For example, if the client node is named AUSTIN, issue the following command:

```
decommission node austin wait=yes
```

- To decommission a virtual machine in the background, issue the DECOMMISSION VM command. For example, if the virtual machine is named AUSTIN, the file space is 7, and the file space name is specified by the file space ID, issue the following command:

```
decommission vm austin 7 nametype=fsid
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid
```

- To decommission a virtual machine in the foreground, issue the DECOMMISSION VM command and specify the `wait=yes` parameter. For example, issue the following command:

```
decommission vm austin 7 nametype=fsid wait=yes
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid wait=yes
```

What to do next

Watch for error messages, which might be displayed in the user interface or in the command output, immediately after you run the process.

You can verify that the client node is decommissioned:

1. On the Operations Center Overview page, click Clients.
 2. In the Clients table, in the At risk column, review the state:
 - o A DECOMMISSIONED state specifies that the node is decommissioned.
 - o A null value specifies that the node is not decommissioned.
 - o A PENDING state specifies that the node is being decommissioned, or the decommission process failed.
- Tip: If you want to determine the status of a pending decommission process, issue the following command:

```
query process
```

3. Review the command output:
 - o If status is provided for the decommission process, the process is in progress. For example:

```
query process
```

Process Number	Process Description	Process Status
3	DECOMMISSION NODE	Number of backup objects deactivated for node NODE1: 8 objects deactivated.

- o If no status is provided for the decommission process, and you did not receive an error message, the process is incomplete. A process can be incomplete if files that are associated with the node are not yet deactivated. After the files are deactivated, run the decommission process again.
- o If no status is provided for the decommission process, and you receive an error message, the process failed. Run the decommission process again.

Related reference:

- [DECOMMISSION NODE \(Decommission a client node\)](#)
- [DECOMMISSION VM \(Decommission a virtual machine\)](#)
- [QUERY NODE \(Query nodes\)](#)
- [REMOVE REPLNODE \(Remove a client node from replication\)](#)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Spectrum Protect™ server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

About this task

Some application clients always save data to the server as active backup data. Because active backup data is not managed by inventory expiration policies, the data is not deleted automatically, and uses server storage space indefinitely. To free the storage space that is used by obsolete data, you can deactivate the data.

When you run the deactivation process, all active backup data that was stored before the specified date becomes inactive. The data is deleted as it expires and cannot be restored. The deactivation feature applies only to application clients that protect Oracle databases.

Procedure

1. From the Operations Center Overview page, click Clients.
2. In the Clients table, select one or more clients and click More > Clean Up.
Command-line method: Deactivate data by using the DEACTIVATE DATA command.

Related reference:

- [DEACTIVATE DATA \(Deactivate data for a client node\)](#)

Managing data storage

Manage your data for efficiency and add supported devices and media to the server to store client data.

- Auditing a storage pool container
Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.

- **Managing inventory capacity**
Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.
- **Managing memory and processor usage**
Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.
- **Tuning scheduled activities**
Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Related reference:

[Storage pool types](#)

Auditing a storage pool container

Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.

About this task

You audit a storage pool container in the following situations:

- When you issue the QUERY DAMAGED command and a problem is detected
- When the server displays messages about damaged data extents
- Your hardware reports an issue and error messages that are associated with the storage pool container are displayed

Procedure

1. To audit a storage pool container, issue the AUDIT CONTAINER command. For example, issue the following command to audit a container, 00000000000076c.dcf:

```
audit container c:\tsm-storage\07\00000000000076c.dcf
```

2. Review the output from the ANR4891I message for information about any damaged data extents.

What to do next

If you detect problems with the storage pool container, you can restore data based on your configuration. You can repair the contents in the storage pool by using the REPAIR STGPOOL command.

Restriction: You can repair the contents of the storage pool only if you protected the storage pool by using the PROTECT STGPOOL command.

Related reference:

[AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

[QUERY DAMAGED](#) (Query damaged data in a directory-container or cloud-container storage pool)

Managing inventory capacity

Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.

Before you begin

The active and archive logs have the following characteristics:

- The active log can be a maximum size of 512 GB. For more information about sizing the active log for your system, see [Planning the storage arrays](#).
- The archive log size is limited to the size of the file system that it is installed on. The archive log size is not maintained at a predefined size like the active log. Archive log files are automatically deleted after they are no longer needed.

As a best practice, you can optionally create an archive failover log to store archive log files when the archive log directory is full.

Check the Operations Center to determine the component of the inventory that is full. Ensure that you stop the server before you increase the size of one of the inventory components.

Procedure

- To increase the size of the database, complete the following steps:
 - Create one or more directories for the database on separate drives or file systems.
 - Issue the EXTEND DBSPACE command to add the directory or directories to the database. The directories must be accessible to the instance user ID of the database manager. By default, data is redistributed across all database directories and space is reclaimed.
- Tips:
- The time that is needed to complete redistribution of data and reclaiming of space is variable, depending on the size of your database. Make sure that you plan adequately.
 - Ensure that the directories that you specify are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.
- Halt and restart the server to fully use the new directories.
 - Reorganize the database if necessary. Index and table reorganization for the server database can help to avoid unexpected database growth and performance issues. For more information about reorganizing the database, see technote 1683633.

- To decrease the size of the database for V7.1 servers and later, issue the following DB2® commands from the server instance directory:

Restriction: The commands can increase I/O activity, and might affect server performance. To minimize performance problems, wait until one command is completed before you issue the next command. The DB2 commands can be issued when the server is running.

```
db2 connect to tsmdb1
db2 set schema tsmdb1
db2 ALTER TABLESPACE USERSPACE1 REDUCE MAX
db2 ALTER TABLESPACE IDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGEIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGESPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE5 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE5 REDUCE MAX
```

- To increase or decrease the size of the active log, complete the following steps:
 1. Ensure that the location for the active log has enough space for the increased log size. If a log mirror exists, its location must also have enough space for the increased log size.
 2. Halt the server.
 3. In the dsmserv.opt file, update the ACTIVELOGSIZE option to the new size of the active log, in megabytes. The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
257 GB - 512 GB	20480 MB

To change the active log to its maximum size of 512 GB, enter the following server option:

```
activelogsiz 524288
```

4. If you plan to use a new active log directory, update the directory name that is specified in the ACTIVELOGDIRECTORY server option. The new directory must be empty and must be accessible to the user ID of the database manager.
 5. Restart the server.
- Compress the archive logs to reduce the amount of space that is required for storage. Enable dynamic compression of the archive log by issuing the following command:

```
setopt archlogcompress yes
```

Restriction: Use caution when you enable the ARCHLOGCOMPRESS server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log directory file system usage nears out of space conditions, the ARCHLOGCOMPRESS server option must be disabled. You can use the SETOPT command to disable archive log compression immediately without halting the server.

Related reference:

- 🔗 [ACTIVELOGSIZE server option](#)
- 🔗 [EXTEND DBSPACE \(Increase space for the database\)](#)
- 🔗 [SETOPT \(Set a server option for dynamic update\)](#)

Managing memory and processor usage

Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.

Before you begin

- Ensure that your configuration uses the required hardware and software. For more information, see IBM Spectrum Protect™ Supported Operating Systems.
- For more information about managing resources such as the database and recovery log, see Planning the storage arrays.
- Add more system memory to determine whether there is a performance improvement. Monitor memory usage regularly to determine whether more memory is required.

Procedure

1. Release memory from the file system cache where possible.
2. To manage the system memory that is used by each server on a system, use the DBMEMPERCENT server option. Limit the percentage of system memory that can be used by the database manager of each server. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.
3. Set the user data limit and private memory for the database to ensure that private memory is not exhausted. Exhausting private memory can result in errors, less than optimal performance, and instability.

Tuning scheduled activities

Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Procedure

1. Monitor system performance regularly to ensure that client backup and server maintenance tasks are completing successfully. Follow the instructions in Monitoring a multisite disk solution.

2. Optional: If the monitoring information shows that the server workload increased, review the planning information. Review whether the capacity of the system is adequate in the following cases:
 - o The number of clients increases
 - o The amount of data that is being backed up increases
 - o The amount of time that is available for backups changes
3. Determine whether your solution is performing at the level you expect. Review the client schedules to check whether tasks are completing within the scheduled time frame:
 - a. On the Clients page of the Operations Center, select the client.
 - b. Click Details.
 - c. From the client Summary page, review the Backed up and Replicated activity to identify any risks. Adjust the time and frequency of client backup operations, if necessary.
4. Schedule enough time for the following maintenance tasks to complete successfully within a 24-hour period:
 - a. Protect storage pools.
 - b. Replicate node data.
 - c. Back up the database.
 - d. Run expiration processing to remove client backups and archive file copies from server storage.

Tip: Schedule maintenance tasks to start at an appropriate time and in the correct sequence. For example, schedule replication tasks after client backups complete successfully.

 - Moving clients from one server to another
To avoid running out of space on a server or to resolve workload issues, you might have to move client nodes from one server to another.

Related concepts:

[Performance](#)

Related tasks:

Defining schedules for server maintenance activities

[Deduplicating data \(V7.1.1\)](#)

Managing replication

Use replication to recover data at a disaster recovery site and to maintain the same level of files on the source and target servers. You can manage replication at the node level. You can also protect data at the storage-pool level.

- Replication compatibility
Before you set up replication operations with IBM Spectrum Protect, you must ensure that the source and target replication servers are compatible for replication.
- Enabling node replication
You can enable node replication to protect your data.
- Protecting data in directory-container storage pools
Protect data in directory-container storage pools to reduce node replication time and to enable repair of data in directory-container storage pools.
- Modifying replication settings
Modify replication settings in the Operations Center. Change settings such as the number of replication sessions, replication rules, the data that you want to replicate, the replication schedule, and the replication workload.
- Setting different retention policies for the source server and target server
You can set policies on the target replication server that manage the replicated client-node data differently than on the source server. For example, you can maintain a different number of versions of files on the source and the target servers.

Replication compatibility

Before you set up replication operations with IBM Spectrum Protect™, you must ensure that the source and target replication servers are compatible for replication.

Table 1. Replication compatibility of server versions

Source replication server version	Compatible versions for the target replication server
V7.1	V7.1 or later
V7.1.1	V7.1 or later
V7.1.3	V7.1.3 or later

Source replication server version	Compatible versions for the target replication server
V7.1.4	V7.1.3 or later
V7.1.5	V7.1.3 or later
V7.1.6	V7.1.3 or later
V7.1.7	V7.1.3 or later
V7.1.8	V7.1.3 or later
V8.1	V7.1.3 or later
V8.1.1	V7.1.3 or later
V8.1.2	V7.1.3 or later
V8.1.3	V7.1.3 or later
V8.1.4	V7.1.3 or later
V8.1.5	V7.1.3 or later

Enabling node replication

You can enable node replication to protect your data.

Before you begin

Ensure that the source and target servers are compatible for replication.

About this task

Replicate the client node to replicate all client data, including metadata. By default, node replication is disabled when you start the server for the first time.

Tips:

- To reduce replication processing time, protect the storage pool before you replicate client nodes. When node replication is started, the data extents that are already replicated through storage pool protection are skipped.
- Replication requires increased amounts of memory and sufficient bandwidth to complete processing. Size the database and its logs to ensure that transactions can complete.

Procedure

To enable node replication, complete the following steps in the Operations Center:

- On the Servers page, click Details.
- On the Details page, click Properties.
- In the Replication section, select Enabled in the Outbound replication field.
- Click Save.

What to do next

Complete the following actions:

- To verify that replication was successful, review the Daily monitoring checklist.
- Linux** If the IBM Spectrum Protect server replicates nodes to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Follow the instructions in Determining whether Aspera FASP technology can optimize data transfer in your system environment.

Related reference:

Replication compatibility

Protecting data in directory-container storage pools

Protect data in directory-container storage pools to reduce node replication time and to enable repair of data in directory-container storage pools.

Before you begin

Ensure that at least one directory-container storage pool exists on the target replication server. When you enable replication in the Operations Center, you can schedule storage pool protection. To configure replication and enable storage pool protection, complete the following steps:

1. On the Operations Center menu bar, hover over Storage and click Replication.
2. On the Replication page, click Server Pair.
3. Complete the steps in the Add Server Pair wizard.

About this task

Protecting a directory-container storage pool backs up data extents to another storage pool, and can improve performance for node replication. When node replication is started, the data extents that are already backed up through storage pool protection are skipped, which reduces the replication processing time. You can schedule the protection of storage pools several times a day to keep up with changes to data.

By protecting a storage pool, you do not use resources that replicate existing data and metadata, which improves server performance. You must use directory-container storage pools if you want to protect and back up the storage pool only.

Alternative protection strategy: As an alternative to using replication, you can protect data in directory-container storage pools by copying the data to container-copy storage pools. Data in container-copy storage pools is stored on tape volumes. Tape copies that are stored offsite provide additional disaster recovery protection in a replicated environment.

Procedure

1. Alternatively, to enable storage pool protection, you can use the PROTECT STGPOOL command from the source server to back up data extents in a directory-container storage pool. For example, to protect a directory-container storage pool that is named POOL1 issue the following command:

```
protect stgpool pool1
```

As part of the operation of the PROTECT STGPOOL command, damaged extents in the target storage pool are repaired. To be repaired, extents must already be marked as damaged on the target server. For example, an AUDIT CONTAINER command might identify damage in the target storage pool before the PROTECT STGPOOL command is issued.

2. Optional: If damaged extents were repaired in the target storage pool and you protect multiple source storage pools in one target storage pool, complete the following steps to ensure a complete repair:
 - a. Issue the PROTECT STGPOOL command for all source storage pools to repair as much of the damage as possible.
 - b. Issue the PROTECT STGPOOL command again for all source storage pools. For this second operation, use the FORCERECONCILE=YES parameter. This step ensures that any repairs from other source pools are properly recognized for all source storage pools.

Results

If a directory-container storage pool is protected, you can repair the storage pool if damage occurs, by using the REPAIR STGPOOL command.

Restriction: If you replicate client nodes but do not protect the directory-container storage pool, you cannot repair the storage pool.

What to do next

Complete the following actions:

1. To view replication workload status, follow the instructions in the Daily monitoring checklist.
2. **Linux** If the IBM Spectrum Protect server replicates nodes to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Follow the instructions in Determining whether Aspera FASP technology can optimize data transfer in your system environment.

Related reference:

[🔗 Repairing and recovering data in directory-container storage pools](#)

🔗 [AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

🔗 [PROTECT STGPOOL](#) (Protect storage pool data)

Related information:

🔗 [Directory-container storage pools FAQs](#)

🔗 [Cloud-container storage pools FAQs](#)

Modifying replication settings

Modify replication settings in the Operations Center. Change settings such as the number of replication sessions, replication rules, the data that you want to replicate, the replication schedule, and the replication workload.

About this task

You might need to customize your replication settings in the following scenarios:

- Changes to data priorities
- Changes to replication rules
- Requirement for a different server to be the target server
- Scheduled processes that negatively affect server performance

Procedure

Use the Operations Center to modify replication settings.

Task	Procedure
Change a replication rule.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.In the Replication section, choose the replication rule that you want to apply: Default archive rule, Default backup rule, or Default space-management rule.Click Save.
Specify the duration that replication records are retained.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.In the Replication section, enter the number of days that replication records must be retained in the Retain replication history field. Alternatively, select the Do not retain check box if you do not require replication records.Click Save.
Specify a target replication server.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.In the Replication section, specify the target server.Click Save.
Cancel a replication process.	<ol style="list-style-type: none">On the Servers page, click Active tasks.Select the process or session that you want to cancel.Click Cancel.

Setting different retention policies for the source server and target server

You can set policies on the target replication server that manage the replicated client-node data differently than on the source server. For example, you can maintain a different number of versions of files on the source and the target servers.

Procedure

1. From the source replication server, validate the replication configuration and verify that the source replication server can communicate with the target replication server by issuing the `VALIDATE REPLICATION` command. For example, validate the configuration by using the name of one client node that is being replicated:

```
validate replication node1 verifyconnection=yes
```

2. From the source replication server, issue the `VALIDATE REPLPOLICY` command to review the differences between the policies on the source and target replication servers. For example, to display the differences between the policies on the source server and the target server, `CVT_SRV2`, issue the following command from the source server:

```
validate replpolicy cvt_srv2
```

3. Update the policies on the target server if necessary.

Tip: You can use the Operations Center to modify the policies on the target server. Follow the instructions in [Editing policies](#).

For example, to maintain inactive versions of files for a shorter time on the target server than on the source server, reduce the Backups setting in the management classes that apply to replicated client data.

4. Enable the target replication server to use its policies to manage the replicated client-node data by issuing the `SET DISSIMILARPOLICIES` command on the source server. For example, to enable the policies on the target replication server, `CVT_SRV2`, issue the following command on the source server:

```
set dissimilarpolicies cvt_srv2 on
```

The next time that the replication process runs, the policies on the target replication server are used to manage the replicated client-node data.

Tip: If you configure replication by using the Operations Center and the policies on the source and target replication servers do not match, the policy that is specified for the source replication server is used. If you enabled the policies on the target replication server by using the `SET DISSIMILARPOLICIES` command, the policy that is specified for the target replication server is used. If the target replication server does not have the policy that is used by the node on the source replication server, the `STANDARD` policy is used.

Related reference:

- [EXPORT POLICY](#) (Export policy information)
- [SET DISSIMILARPOLICIES](#) (Enable the policies on the target replication server to manage replicated data)
- [VALIDATE REPLICATION](#) (Validate replication for a client node)
- [VALIDATE REPLPOLICY](#) (Verify the policies on the target replication server)

Securing the server

Secure the IBM Spectrum Protect™ server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.

- Security concepts
You can protect IBM Spectrum Protect from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.
- Managing administrators
An administrator who has system authority can complete any task with the IBM Spectrum Protect server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.
- Changing password requirements
You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect.
- Securing IBM Spectrum Protect on the system
Protect the system where the IBM Spectrum Protect server runs to prevent unauthorized access.

Security concepts

You can protect IBM Spectrum Protect™ from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.

Transport Layer Security

You can use the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol to provide transport layer security for a secure connection between servers, clients, and storage agents. If you send data between the server, client, and storage agent, use SSL or TLS to encrypt the data.

Tip: Any IBM Spectrum Protect documentation that indicates "SSL" or to "select SSL" applies to TLS.

SSL is provided by the Global Security Kit (GSKit) that is installed with the IBM Spectrum Protect server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with a DB2® database instance that is used by any IBM Spectrum Protect servers.

Each server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a certificate authority (CA). You can use your own certificates or purchase certificates from a CA. Either certificate must be installed and added to the key database on the IBM Spectrum Protect server, client, or storage agent. The certificate is verified by the SSL client or server that requests or initiates the SSL communication. Some CA certificates are preinstalled in the key databases, by default.

SSL is set up independently on the IBM Spectrum Protect server, client, and storage agent.

Authority levels

With each IBM Spectrum Protect server, different administrative authority levels are available that determine the tasks that an administrator can complete.

After registration, an administrator must be granted authority by being assigned one or more administrative authority levels. An administrator with system authority can complete any task with the server and assign authority levels to other administrators by using the GRANT AUTHORITY command. Administrators with policy, storage, or operator authority can complete subsets of tasks.

An administrator can register other administrator IDs, grant levels of authority to them, rename IDs, remove IDs, and lock and unlock them from the server.

An administrator can control access to specific client nodes for root user IDs and non-root user IDs. By default, a non-root user ID cannot back up data on the node. Use the UPDATE NODE command to change the node settings to enable backup.

Passwords

By default, the server automatically uses password authentication. With password authentication, all users must enter a password when they access the server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see Managing passwords and logon procedures (V7.1.1).

Table 1. Password authentication characteristics

Characteristic	More information
Case-sensitivity	Not case-sensitive.
Default password expiration	90 days. The expiration period begins when an administrator ID or client node is first registered to the server. If the password is not changed within this period, the password must be changed the next time that the user accesses the server.
Invalid password attempts	You can set a limit on consecutive invalid password attempts for all client nodes. When the limit is exceeded, the server locks the node.
Default password length	8 characters. The administrator can specify a minimum length. Beginning with Version 8.1.4, the default minimum length for server passwords changed from 0 to 8 characters.

Session security

Session security is the level of security that is used for communication among IBM Spectrum Protect client nodes, administrative clients, and servers and is set by using the SESSIONSECURITY parameter.

The SESSIONSECURITY parameter can be set to one of the following values:

- The STRICT value enforces the highest level of security for communication between IBM Spectrum Protect servers, nodes, and administrators.
- The TRANSITIONAL value specifies that the existing communication protocol is used while you update your IBM Spectrum Protect software to V8.1.2 or later. This is the default. When SESSIONSECURITY=TRANSITIONAL, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

Note: You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Spectrum Protect servers, but you are not required to upgrade servers first. Communication between servers and clients is not interrupted.

For more information about the SESSIONSECURITY parameter values, see the following commands.

Table 2. Commands used to set the SESSIONSECURITY parameter

Entity	Command
Client nodes	<ul style="list-style-type: none">• REGISTER NODE• UPDATE NODE
Administrators	<ul style="list-style-type: none">• REGISTER ADMIN• UPDATE ADMIN
Servers	<ul style="list-style-type: none">• DEFINE SERVER• UPDATE SERVER

Administrators that authenticate by using the DSMADMC command, DSMC command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

Tips:

- Ensure that all IBM Spectrum Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.
- After an administrator authenticates to a V8.1.2 or later server by using a V8.1.2 or later client, the administrator can authenticate only on clients or servers that are using V8.1.2 or later. An administrator command can be issued from any system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Enforce the highest level of security for communication with the IBM Spectrum Protect server by ensuring that all nodes, administrators, and servers use STRICT session security. You can use the SELECT command to determine which servers, nodes, and administrators are using TRANSITIONAL session security and should be updated to use STRICT session security.

Related tasks:

[↗ Securing communications](#)

Managing administrators

An administrator who has system authority can complete any task with the IBM Spectrum Protect™ server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.

Procedure

Complete the following tasks to modify administrator settings.

Task	Procedure
Add an administrator.	<p>To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:</p> <ol style="list-style-type: none"> Register the administrator and specify Pa\$#\$twO as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$twO</pre> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>
Change administrative authority.	<p>Change the authority level for an administrator, ADMIN1.</p> <ul style="list-style-type: none"> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre> Revoke system authority for the administrator by issuing the following command: <pre>revoke authority admin1 classes=system</pre>
Remove administrators.	<p>Remove an administrator, ADMIN1, from accessing the IBM Spectrum Protect server by issuing the following command:</p> <pre>remove admin admin1</pre>
Temporarily prevent access to the server.	<p>Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.</p>

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect™.

About this task

By enforcing password authentication and managing password restrictions, you protect your data and your servers from potential security risks.

Procedure

Complete the following tasks to change password requirements for IBM Spectrum Protect servers.

Table 1. Authentication tasks for IBM Spectrum Protect servers

Task	Procedure
Set a limit for invalid password attempts.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details, and then click the Properties tab. Set the number of invalid attempts in the Invalid sign-on attempt limit field. <p>The default value at installation is 0.</p>

Task	Procedure
Set a minimum length for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of characters in the Minimum password length field.
Set the expiration period for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of days in the Password common expiration field.
Disable password authentication.	<p>By default, the server automatically uses password authentication. With password authentication, all users must enter a password to access the server.</p> <p>You can disable password authentication only for passwords that authenticate with the server (LOCAL). By disabling password authentication, you increase the security risk for the server.</p>
Set a default authentication method.	<p>Issue the SET DEFAULTAUTHENTICATION command. For example, to use the server as the default authentication method, issue the following command:</p> <pre>set defaultauthentication local</pre> <p>To update one client node to authenticate with the server, include AUTHENTICATION=LOCAL in the UPDATE NODE command:</p> <pre>update node authentication=local</pre>

Related concepts:

- [Authenticating IBM Spectrum Protect users by using an LDAP server](#)
- [Managing passwords and logon procedures \(V7.1.1\)](#)

Securing IBM Spectrum Protect on the system

Protect the system where the IBM Spectrum Protect™ server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the directories for the server database and the server instance. Keep the access settings for these directories that you configured during implementation.

- Restricting user access to the server
Authority levels determine what an administrator can do with the IBM Spectrum Protect server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.
- Limiting access through port restrictions
Limit access to the server by applying port restrictions.

Restricting user access to the server

Authority levels determine what an administrator can do with the IBM Spectrum Protect™ server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.

Procedure

1. After you register an administrator by using the REGISTER ADMIN command, use the GRANT AUTHORITY command to set the administrator's authority level. For details about setting and changing authority, see Managing administrators.
2. To control the authority of an administrator to complete some tasks, use the following two server options:
 - a. You can select the authority level that an administrator must have to issue QUERY and SELECT commands with the QUERYAUTH server option. By default, no authority level is required. You can change the requirement to one of the authority levels, including system.
 - b. You can specify that system authority is required for commands that cause the server to write to an external file with the REQSYSAUTHOUTFILE server option. By default, system authority is required for such commands.
3. You can restrict data backup on a client node to only root user IDs or authorized users. For example, to limit backups to the root user ID, issue the REGISTER NODE or UPDATE NODE command and specify the BACKUPINITIATION=root parameter:

```
update node backupinitiation=root
```

Limiting access through port restrictions

Limit access to the server by applying port restrictions.

About this task

You might have to restrict access to specific servers, based on your security requirements. The IBM Spectrum Protect™ server can be configured to listen on four TCP/IP ports: two that can be used for either regular TCP/IP protocols or Secure Sockets Layer (SSL)/Transport Layer Security (TLS) protocols and two that can be used only for the SSL/TLS protocol.

Procedure

You can set the server options to specify the port that you require, as listed in Table 1.

Table 1. Server options and port access

Server option	Port access
TCPPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default value is 1500.
TCPADMINPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default is the value of TCPPORT. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.
SSLTCPPORT	Specifies the SSL TCP/IP port address for a server. This port listens for SSL-enabled sessions only. A default port value is not available.
SSLTCPADMINPORT	Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. A default port value is not available. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (SSLTCPPORT and SSLTCPADMINPORT):

- When you specify the server's SSL-only port for the LLADDRESS on the DEFINE SERVER or UPDATE SERVER command, you must also specify the SSL=YES parameter.
- When you specify the server's SSL-only port for the client's TCPPORT option, you must also specify YES for the SSL client option.

Related reference:

Planning firewall access

Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.

Before you begin

You must have system or operator privilege to stop and start the IBM Spectrum Protect™ server.

- Stopping the server
Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.
- Starting the server for maintenance or reconfiguration tasks
Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Stopping the server

Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.

About this task

When you issue the HALT command to stop the server, the following actions occur:

- All processes and client node sessions are canceled.
- All current transactions are stopped. (The transactions will be rolled back when the server is restarted.)

Procedure

To prepare the system and stop the server, complete the following steps:

1. Prevent new client node sessions from starting by issuing the DISABLE SESSIONS command:

```
disable sessions all
```

2. Determine whether any client node sessions or processes are in progress by completing the following steps:
 - a. On the Overview page of the Operations Center, view the Activity area for the total numbers of processes and sessions that are currently active. If numbers differ significantly from the usual numbers that are displayed during your daily storage-management routine, view other status indicators in the Operations Center to check whether there is a problem.
 - b. View the graph in the Activity area to compare the amount of network traffic over the following periods:
 - The current period, that is, the most recent 24-hour period
 - The previous period, that is, the 24 hours before the current periodIf the graph for the previous period represents the expected amount of traffic, significant differences on the graph for the current period might indicate a problem.
 - c. On the Servers page, select a server for which you want to view processes and sessions, and click Details. If the server is not registered as a hub or spoke server in the Operations Center, obtain information about processes by using administrative commands. Issue the QUERY PROCESS command to query processes and obtain information about sessions by issuing the QUERY SESSION command.
3. Wait until the client node sessions are completed or cancel them. To cancel processes and sessions, complete the following steps:
 - On the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - Click Cancel.
 - If the server is not registered as a hub or spoke server in the Operations Center, cancel sessions by using administrative commands. Issue the CANCEL SESSION command to cancel a session and cancel processes by using the CANCEL PROCESS command.
Tip: If the process that you want to cancel is waiting for a tape volume to be mounted, the mount request is canceled. For example, if you issue an EXPORT, IMPORT, or MOVE DATA command, the command might initiate a process that requires a tape volume to be mounted. However, if a tape volume is being mounted by an automated library, the cancel operation might not take effect until the mount process is complete. Depending on your system environment, this could take several minutes.
4. Stop the server by issuing the HALT command:

```
halt
```

Starting the server for maintenance or reconfiguration tasks

Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSErv utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmserve.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserve maintenance
```




Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode. Follow the instructions for your operating system:
 - o  Starting the server instance
 - o  Starting the server instance
 - o  Starting the server instance

Operations that were disabled during maintenance mode are reenabled.

Planning to upgrade the server

When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect™ server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.

About this task

Follow these guidelines:

- The preferred method is to upgrade the server by using the installation wizard. After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.
- If upgrades are available for both the server component and the Operations Center component, select the check boxes to upgrade both components.

Procedure

1. Review the list of fix packs and interim fixes. See technote 1239415.
2. Review product improvements, which are described in readme files.
Tip: When you obtain the installation package file from the IBM Spectrum Protect support site, you can also access the readme file.
3. Ensure that the version that you upgrade your server to is compatible with other components, such as storage agents and library clients. See technote 1302789.
4. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted. See technote 1053218.
5. Review the upgrade instructions. Ensure that you back up the server database, the device configuration information, and the volume history file.

What to do next

To install a fix pack or interim fix, follow the instructions for your operating system:

- **AIX** Installing an IBM Spectrum Protect server fix pack
- **Linux** Installing an IBM Spectrum Protect server fix pack
- **Windows** Installing an IBM Spectrum Protect server fix pack

Related information:

[Upgrade and Migration Process - Frequently Asked Questions](#)

Preparing for an outage or system update

Prepare IBM Spectrum Protect™ to maintain your system in a consistent state during a planned power outage or system update.

About this task

Ensure that you schedule activities regularly to manage, protect, and maintain the server.

Procedure

1. Cancel processes and sessions that are in progress by completing the following steps:
 - a. In the Operations Center, on the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - b. Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - c. Click Cancel.
2. Stop the server by issuing the HALT command:

```
halt
```

Tip: You can issue the halt command from the Operations Center by hovering over the Settings icon and clicking Command Builder. Then, select the server, type `halt`, and press Enter.

Implementing a disaster recovery plan

Implement a disaster recovery strategy to recover your applications if a disaster occurs and to ensure high server availability.

About this task

Determine your disaster recovery requirements by identifying the business priorities for client node recovery, the systems that you use to recover data, and whether client nodes have connectivity to a recovery server. Use replication and storage pool protection to protect data. You must also determine how often directory-container storage pools are protected.

- **Completing recovery drills**
Schedule disaster recovery drills to prepare for audits that certify the recoverability of the IBM Spectrum Protect server and to ensure that data can be restored and operations can resume after an outage. A drill also helps you ensure that all data can be restored and operations resumed before a critical situation occurs.

Recovering from data loss or system outages

You can use IBM Spectrum Protect™ to recover data that was lost when a disaster or system outage occurred. You can recover directory-container storage pools, client data, and databases.

Before you begin

Schedule client and server workloads to achieve the best performance for your storage environment. Issue the PROTECT STGPOOL and REPLICATE NODE commands as part of the schedule. Protect the storage pool before you replicate the client node. When node replication is started, the data extents that are already replicated through storage pool protection are skipped, which reduces replication processing time.

Procedure

Use the following recovery methods based on the component that you must recover.

Component to recover	Procedure	More information
Directory-container storage pool	<p>To recover directory-container storage pools, complete the following steps:</p> <ol style="list-style-type: none">a. Scan for damaged data extents in the directory-container storage pool by using the AUDIT CONTAINER command and specifying the ACTION=SCANALL parameter.b. Repair damaged data extents in the directory-container storage pool by using the REPAIR STGPOOL command. Restriction: You can repair a storage pool only if the storage pool is protected.c. Remove damaged data extents by using the AUDIT CONTAINER command and specifying the ACTION=REMOVEDAMAGED parameter.	Repairing storage pools

Component to recover	Procedure	More information
Client data	<p>Prerequisites:</p> <ul style="list-style-type: none"> The source replication server, the target replication server, and the client must be at the V7.1 level or later. If any of the servers are at an earlier level, automatic failover is disabled and you must rely on manual failover. <p>Manually configure the client to automatically fail over to the target server for data recovery.</p> <p>If you enabled the client for automated client failover, you can recover the data by using automatic failover function. You can verify that the <code>userreplicationfailover</code> option is either not in the client options file or is set to <code>yes</code>. Recover data from the target server when the source server is unavailable due to an outage by using automatic failover.</p> <p>Tip:</p> <ul style="list-style-type: none"> Use the <code>SET FAILOVERHLADDRESS</code> command to specify the IP address for the replication server during failover, if the address is different from the IP address that is specified for the replication process. 	<ul style="list-style-type: none"> Recovering damaged data from a replicated copy <code>SET FAILOVERHLADDRESS</code> (Set a failover high level address)
Database	<p>Prerequisites:</p> <ul style="list-style-type: none"> To restore the database after a disaster, you must have a copy of the current device configuration file. The device configuration file cannot be recreated. Ensure that you have a backed up version of the database. <p>Restore the IBM Spectrum Protect database to the most current state or to a specific point in time by using the <code>DSMSERV RESTORE DB</code> server utility.</p>	<p><code>DSMSERV RESTORE DB</code> (Restore the database)</p>

- Restoring the database
You might have to restore the IBM Spectrum Protect database after a disaster. You can restore the database to the most current state or to a specified point in time. You must have full, incremental, or snapshot database backup volumes to restore the database.
- Recovering damaged data from a replicated copy
If a source replication server is unavailable, you can recover damaged data from a replicated copy that is stored on the

- target replication server.
 - Repairing storage pools
- If a disaster or system outage occurred, you can repair deduplicated data extents in a directory-container storage pool.

Related reference:

- [AUDIT CONTAINER \(Verify the consistency of database information for a directory-container storage pool\)](#)
- [DSMSERV RESTORE DB \(Restore the database\)](#)

Restoring the database

You might have to restore the IBM Spectrum Protect™ database after a disaster. You can restore the database to the most current state or to a specified point in time. You must have full, incremental, or snapshot database backup volumes to restore the database.

Before you begin

If the database and recovery log directories are lost, re-create them before you issue the DSMSERV RESTORE DB server utility. For example, use the following commands:

```
mkdir /tsmdb001
mkdir /tsmdb002
mkdir /tsmdb003
mkdir /activelog
mkdir /archlog
mkdir /archfaillog
```

Windows

```
mkdir e:\tsm\db001
mkdir f:\tsm\db001
mkdir g:\tsm\db001
mkdir h:\tsm\activelog
mkdir i:\tsm\archlog
mkdir j:\tsm\archfaillog
```

Restrictions:

- To restore the database to its latest version, you must locate the archive log directory. If you are cannot locate the directory, you can restore the database only to a point in time.
- You cannot use Secure Sockets Layer (SSL) for database restore operations.
- If the release level of the database backup is different from the release level of the server that is being restored, you cannot restore the server database. For example, if you are using a Version 8.1 server and you try to restore a Version 7.1 database, an error occurs.

About this task

Point-in-time restore operations are typically used for situations such as disaster recovery or to remove the effects of errors that can cause inconsistencies in the database. To recover the database to the time when the database was lost, recover the database to its latest version.

Procedure

Use the DSMSERV RESTORE DB server utility to restore the database. Depending on the version of the database that you want to restore, choose one of the following methods:

- Restore a database to its latest version. For example, use the following command:

```
dsmserv restore db
```

- Restore a database to a point in time. For example, to restore the database to a backup series that was created on 19 April 2015, use the following command:

```
dsmserv restore db todate=04/19/2015
```

What to do next

If you restored the database and directory-container storage pools exist on the server, you must identify inconsistencies between the database and the file system.

1. If you restored the database to a point in time and you did not delay reuse of the directory-container storage pool, you must audit all the containers. To audit all containers, issue the following command:

```
audit container stgpool
```

2. If the server cannot identify containers on the system, complete the following steps to display a list of containers:
 - a. From an administrative client, issue the following command:

```
select container_name from containers
```

- b. From the file system, issue the following command for the storage pool directory on the source server:

Tip: The storage pool directory is displayed in the command output:

AIX | **Linux**

```
[root@source]$ ls -lR
```

Windows

```
c:\source_stgpooldir>dir /s
```

- c. Compare the containers that are listed on the file system and the server.
- d. Issue the AUDIT CONTAINER command and specify the container that is missing from the server output. Specify the ACTION=REMOVEDAMAGED parameter to delete the container.
- e. To ensure that the containers are deleted on the file system, review the messages that are displayed.
Tip: The IBM Spectrum Protect server does not recognize containers that are created after the last database backup. Delete extra files that exist on your local file system when compared to the files that exist on the IBM Spectrum Protect server.

Related tasks:

[Replicating client node data after a database restore \(V7.1.1\)](#)

Related reference:

[AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

[DSMSERV RESTORE DB](#) (Restore the database)

Recovering damaged data from a replicated copy

If a source replication server is unavailable, you can recover damaged data from a replicated copy that is stored on the target replication server.

Before you begin

The server name that you specify with the SET REPLSERVER command must match the name of an existing server definition. It must also be the name of the server to be used as the target replication server. If the server name specified by this command does not match the server name of an existing server definition, the command fails.

Tip:

- Use care when you change or remove a target replication server. If you change a target replication server, client-node data that is replicated is sent to a different target replication server. If you remove a target replication server, client node data is not replicated.

Procedure

1. Verify the replication status of the data on the target server. The replication status indicates whether the most recent backup was replicated to the secondary server.
2. Restore data from a target replication server by setting the source replication server as the target replication server. For example, if you want to set the source replication server as the target replication server, server1, issue the following command:

```
set replserver server1
```

What to do next

When you restore the IBM Spectrum Protect™ database on a source replication server, replication is automatically disabled. Before you re-enable replication, determine whether copies of data that are on the target replication server are needed.

Related tasks:

[Replicating client node data after a database restore \(V7.1.1\)](#)

Repairing storage pools

If a disaster or system outage occurred, you can repair deduplicated data extents in a directory-container storage pool.

Before you begin

Identify inconsistencies between the database and the directory-container storage pool by using the `AUDIT CONTAINER` command. By identifying the damaged data extents in the directory-container storage pool, you can determine what data extents to repair.

Before you repair a storage pool, ensure that the storage pool is protected by using the `PROTECT STGPOOL` command.

Procedure

1. To repair a directory-container storage pool, use the `REPAIR STGPOOL` command. For example, to repair a storage pool, `STGPOOL1`, issue the following command:

```
repair stgpool stgpool1
```

2. If the damaged storage pool is specified as a target storage pool on the `PROTECT STGPOOL` command for one or more source storage pools, issue the `PROTECT STGPOOL` command for all source storage pools.
3. To ensure that all damaged data is identified and repaired from other source storage pools, issue the `PROTECT STGPOOL` command again from all source storage pools and specify the `FORCERECONCILE=YES` parameter.
4. To remove objects that refer to damaged data, issue the `AUDIT CONTAINER` command and specify the `ACTION=REMOVEDAMAGED` parameter.
5. If the damaged storage pool is a target storage pool for node replication from one or more source servers, issue the `REPLICATE NODE` command again from all source servers.
6. When the damage is repaired, issue the `PROTECT STGPOOL` command to ensure that the storage pool is protected to another directory-container storage pool.

What to do next

Ensure that no damaged data extents are displayed in the output by using the `QUERY DAMAGED` command.

Related reference:

- [Repairing and recovering data in directory-container storage pools](#)
- [AUDIT CONTAINER \(Verify the consistency of database information for a directory-container storage pool\)](#)
- [QUERY DAMAGED \(Query damaged data in a directory-container or cloud-container storage pool\)](#)
- [REPAIR STGPOOL \(Repair a directory-container storage pool\)](#)

Tape solution

This data protection solution provides storage to tape media, a flexible and affordable option for long-term data retention.

- **Planning for a tape-based data protection solution**
Plan for a data protection solution that includes disk-to-disk-to-tape and disk-to-tape backup operations to optimize storage.
- **Implementation of a tape-based data protection solution**
Implement the tape-based solution, which uses disk-to-disk-to-tape backup and disk staging to optimize storage. By implementing the tape solution, you can enable long-term data retention and achieve low-cost scalability.
- **Monitoring a tape solution**
After you implement an IBM Spectrum Protect tape-based solution, monitor the solution to ensure correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.
- **Managing operations for a tape solution**
Use this information to manage operations for a tape implementation for an IBM Spectrum Protect server.


Planning for a tape-based data protection solution

Plan for a data protection solution that includes disk-to-disk-to-tape and disk-to-tape backup operations to optimize storage.

Planning roadmap

Plan for the tape solution by reviewing the architecture layout in Figure 1 and then completing the roadmap tasks that follow the diagram.

Figure 1. Tape solution



In this data protection configuration, the server uses both disk and tape storage hardware. Storage pool staging is used, in which client data is initially stored in disk storage pools and then later migrated to tape storage pools. For disaster recovery, tape volumes can be stored offsite. Offsite options include physically moving a second copy offsite by a courier or electronically vaulting copies offsite to a remote library.

Tip: The described solution does not include node replication. However, if you want to use node replication to back up a storage pool from disk to disk, ensure that the replication operation is completed before data is migrated from disk to tape. You can also use node replication to back up a storage pool on a local tape device to a copy storage pool on a local tape device.

To plan for a tape-based solution, complete the following tasks:

1. Meet system requirements for hardware and software.
2. Record values for your system configuration in the planning worksheets.
3. Plan for disk storage.
4. Plan for tape storage.
5. Plan for security.

Tape planning requirements

Before you implement a tape solution, review the general guidelines about system requirements. Determine whether to back up data to disk or tape, or a combination of both.

Network bandwidth

The network must have sufficient bandwidth for the expected data transfers between the client and the server, and for the cross-site restore operations that are required for disaster recovery. Use a storage area network (SAN) for data transfers among the server, disk devices, and tape devices. For more information, see Hardware requirements.

Data migration

Migrate all data from disk to tape daily. Specify a FILE device class for disk-based storage pools. Schedule migration to control when processing occurs. To prevent automatic migration based on the migration threshold, specify a value of 100 for the HIGHMIG parameter and 0 for the LOWMIG parameter when you issue the DEFINE STGPOOL command. You must keep at least 20% of the tape drives available for restore operations. To use up to 80% of available tape drives and improve throughput performance, specify the MIGPROCESS parameter.

Consider the following information based on the type of data that is migrated:

- Use tape to back up data from clients that have large objects, such as databases.
Tip: Check with your tape-drive manufacturer for guidance about the size of the database that is suitable to write to tape.
- Use disk to back up data from clients that have smaller objects.
- To back up data directly to tape, use LAN-free data movement. For more information, see [Configuring LAN-free data movement](#).
- Do not back up virtual machines to tape. Use a separate disk-based storage pool that does not migrate to a tape-based storage pool. For more information about virtual machine support, see [technote 1239546](#).

Storage pool capacity

Maintain enough storage pool capacity to allow for 2 days of client backups and a buffer of 20%. You might have to schedule full backups over a few days to ensure that you have enough storage pool space.

Tape drives

Review the manufacturer specifications and estimate the capacity of a tape drive. Determine the amount of space that is required for backup and migration operations. Reserve 20% of tape drives for restore operations.

Related reference:

[MIGRATE STGPOOL \(Migrate storage pool to next storage pool\)](#)

System requirements for a tape-based solution

Hardware and software requirements are provided for a tape-based storage solution that has a data ingestion rate of 14 TB per hour.

Review the information to determine the hardware and software requirements for your storage environment. You might have to make adjustments based on your system size.

- **Hardware requirements**
Hardware requirements for your IBM Spectrum Protect™ solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.
- **Software requirements**
Documentation for the IBM Spectrum Protect tape-based solution includes installation and configuration tasks for IBM® AIX®, Linux, and Microsoft Windows operating systems. You must meet the minimum software requirements that are listed.

Hardware requirements

Hardware requirements for your IBM Spectrum Protect™ solution are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.

For more information about planning disk devices, see [Planning for disk storage](#).

For more information about planning tape devices, see [Planning for tape storage](#).

The following table includes minimum hardware requirements for the server and storage. If you are using local partitions (LPARs) or work partitions (WPARs), adjust the network requirements to take account of the partition sizes. The figures in the table are based on a data ingestion rate of 14 TB per hour.

Hardware component	System requirements
--------------------	---------------------

Hardware component	System requirements
Server processor	<p>AIX 8 processor cores, 3.42 GHz or faster.</p> <p>For example, use a POWER8® processor-based server.</p> <p>Linux Windows 16 processor cores, 2.0 GHz or faster.</p> <p>For example, use an Intel Xeon processor.</p>
Server memory	64 GB RAM.
Network	<p>The following sizing manages approximately 14 TB of data per hour:</p> <ul style="list-style-type: none"> • 10 Gb Ethernet (a minimum of four ports) • 8 Gb Fibre Channel adapter (a minimum of four ports) <p>The number of ports depends on the percentage of daily data ingestion to disk storage pools versus tape storage.</p> <p>Use separate Fibre Channel adapters for tape and disk data.</p>
Storage	<p>Disk</p> <p>Based on the amount of data that you are writing to disk, specify the number of disks that you require.</p> <p>Ensure that the sequential input/output (I/O) throughput of the storage area network (SAN) matches the I/O throughput for the network in the previous row.</p> <p>For example, if you must back up 10 TB of data in a four-hour window, the throughput is approximately 700 MB per second. In this case, the server requires a front-end network (client-to-server path) that supports a minimum throughput of 700 MB per second. The back-end SAN (the server-to-storage device path) also must support a minimum throughput of 700 MB per second.</p> <p>To calculate the required disk speed, use the following formulas:</p> $\frac{(\text{Total amount of daily data ingestion} - \text{amount of daily data ingestion directly to tape})}{(\text{Number of hours for daily client backup operations})} = \text{Megabytes of data ingestion to disk per hour}$ $\frac{(\text{Megabytes of data ingestion to disk per hour})}{(3600 \text{ seconds per hour})} = \text{Megabytes of data ingestion per second that must be supported by the disk technology}$ <p>Tape</p> <p>Select the tape technology that best fits your business requirements. For example, use IBM Linear Tape-Open (LTO) or IBM TS1150 tape drives. Ensure that you have sufficient mount points for client backup operations and for migration. For more information about planning tape storage, see Planning for tape storage. For a list of supported tape devices, see IBM® Support Portal for IBM Spectrum Protect.</p> <p>Tip: To optimize data movement, use LAN-free data movement.</p>
SAN I/O adapters	<p>Segregate disk and tape I/O. For more information about selecting an adapter, see the documentation for Brocade hardware products and for IBM Storwize® storage solutions.</p> <p>Disk</p> <p>Use at least two adapters.</p> <p>Tape</p> <p>Use at least two adapters.</p>

Estimating space requirements for the Operations Center

Hardware requirements for the Operations Center are included in the preceding table, except for the database and archive log space (inventory) that the Operations Center uses to hold records for managed clients.

If you do not plan to install the Operations Center on the same system as the IBM Spectrum Protect server, you can estimate system requirements separately. To calculate system requirements for the Operations Center, see the system requirements calculator in technote 1641684.

Managing the Operations Center on the IBM Spectrum Protect server is a workload that requires extra space for database operations on both the hub server and any spoke servers. The amount of space on the hub server for the archive log is larger if the hub server is monitoring one or more spoke servers. Review the following guidelines to estimate how much space your IBM Spectrum Protect server requires.

Database space for the Operations Center

The Operations Center uses approximately 4.4 GB of database space for every 1000 clients that are monitored on that server. This calculation applies to both hub servers and spoke servers within a configuration.

For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1000 clients. This configuration has a total of 5000 clients across the four servers. Each of the spoke servers requires 4.4 GB of database space. If the spoke servers are at IBM Spectrum Protect Version 8.1.2 or later, the hub server requires 8.8 GB of database space for monitoring only its 2000 clients:

$$(4.4 \text{ GB} \times 2) = 8.8 \text{ GB}$$

Database space for managed data

Managed data is the amount of data that is protected, including the amount of data for all retained versions.

- For client types that perform incremental-forever backups, the following formula can be used to estimate the total managed data:

$$\text{Front-end} + (\text{front-end} \times \text{change rate} \times (\text{retention} - 1))$$

For example, if you back up 100 TB of front-end data, use a 30-day retention period, and have a 5% change rate, calculate your total managed data by using the following figures:

$$100 \text{ TB} + (100 \text{ TB} \times 0.05 \times (30-1)) = 245 \text{ TB total managed data}$$

- For client types that perform full backups every day, the following formula can be used to estimate the total managed data:

$$\text{Front-end} \times \text{retention} \times (1 + \text{change rate})$$

For example, if you back up 10 TB of front-end data, use a 30-day retention period, and have a 3% change rate, calculate your total managed data by using the following figures:

$$10 \text{ TB} \times 30 \times (1 + .03) = 309 \text{ TB total managed data}$$

Unstructured data, average object size: 4 MB

Structured data, average object size: 128 MB

Unstructured data, number of objects =

$$(245 \text{ TB} \times 1024 \times 1024) / 4 \text{ MB} = 64225280$$

Structured data, number of objects =

$$(309 \text{ TB} \times 1024 \times 1024) / 128 \text{ MB} = 2531328$$

Total number of objects: 66756608

Managed data cost (1 KB per object) =

$$(66756608 \text{ KB}) / (1024 \times 1024) = 63.66 \text{ GB}$$

Plan for 20% of additional space so that database systems are not at 100% capacity:

$$\text{Database total physical storage requirements} = (\text{managed data space} + \text{Operations Center space}) \times (1.20)$$

For this example, you would calculate the space by using the following figures:

$$(66.33 \text{ GB} + 8.4 \text{ GB}) \times 1.20 = 76.41 \text{ GB}$$

Archive log space

The Operations Center uses approximately 18 GB of archive log space every 24 hours, per server, for every 1000 clients monitored on that server. Additionally, for every 1000 clients that are monitored on spoke servers, additional archive log space is used on the hub server. For spoke servers at V8.1.2 or later, this added amount is 1.2 GB of archive log space on the hub server per 1000 clients monitored every 24 hours.

For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1000 clients. This configuration has a total of 5000 clients across the four servers. You can calculate the archive log space for the hub server by using the following formula:

$$((18 \text{ GB} \times 2) + (1.2 \text{ GB} \times 3)) = 39.6 \text{ GB of archive log space}$$

These estimates are based on the default status collection interval of 5 minutes. If you reduce the collection interval from once every 5 minutes to once every 3 minutes, the space requirements increase. The following examples show the approximate increase in the log space requirements with a collection interval of once every 3 minutes for a configuration in which V8.1.2 or later spoke servers are monitored:

- Hub server: In the range 39.6 GB - 66 GB
- Each spoke server: In the range 18 GB - 30 GB

Allocate archive log space so that you can support the Operations Center without affecting server operations.

Software requirements

Documentation for the IBM Spectrum Protect™ tape-based solution includes installation and configuration tasks for IBM® AIX®, Linux, and Microsoft Windows operating systems. You must meet the minimum software requirements that are listed.

For information about software requirements for IBM lin_tape device drivers, refer to the IBM Tape Device Drivers Installation and User's Guide.

AIX systems

Type of software	Minimum software requirements
Operating system	IBM AIX 7.1 For more information about operating system requirements, see AIX: Minimum system requirements for AIX systems.
Gunzip utility	The gunzip utility must be available on your system before you install or upgrade the IBM Spectrum Protect server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.
File system type	JFS2 file systems AIX systems can cache a large amount of file system data, which can reduce memory that is required for server and IBM DB2® processes. To avoid paging with the AIX server, use the rbrw mount option for the JFS2 file system. Less memory is used for the file system cache and more is available for IBM Spectrum Protect. Do not use the file system mount options, Concurrent I/O (CIO), and Direct I/O (DIO), for file systems that contain the IBM Spectrum Protect database, logs, or storage pool volumes. These options can cause performance degradation of many server operations. IBM Spectrum Protect and DB2 can still use DIO where it is beneficial to do so, but IBM Spectrum Protect does not require the mount options to selectively take advantage of these techniques.
Other software	Korn Shell (ksh)

Linux systems

Type of software	Minimum software requirements
------------------	-------------------------------

Type of software	Minimum software requirements
Operating system	Red Hat Enterprise Linux 7 (x86_64)
Libraries	GNU C libraries, Version 2.3.3-98.38 or later that is installed on the IBM Spectrum Protect system. Red Hat Enterprise Linux Servers: <ul style="list-style-type: none"> • libaio • libstdc++.so.6 (32-bit and 64-bit packages are required) • numactl.x86_64
File system type	Format database-related file systems with ext3 or ext4. For storage pool-related file systems, use XFS.
Other software	Korn Shell (ksh)

Windows systems

Type of software	Minimum software requirements
Operating system	Microsoft Windows Server 2012 R2 (64-bit) or Windows Server 2016
File system type	NTFS
Other software	Windows 2012 R2 or Windows 2016 with .NET Framework 3.5 is installed and enabled. The following User Account Control policies must be disabled: <ul style="list-style-type: none"> • User Account Control: Admin Approval Mode for the Built-in Administrator account • User Account Control: Run all administrators in Admin Approval Mode

Planning worksheets

Use the planning worksheets to record values that you use to set up your system and configure the IBM Spectrum Protect™ server. Use the best practice default values that are listed in the worksheets.

Each worksheet helps you prepare for different parts of the system configuration by using best practice values:

Server system preconfiguration

Use the preconfiguration worksheets to plan for the file systems and directories that you create when you configure file systems for IBM Spectrum Protect during system setup. All directories that you create for the server must be empty.

Server configuration

Use the configuration worksheets when you configure the server. Default values are suggested for most items, except where noted.

Table 1. Worksheet for preconfiguration of a server system

Item	Default value	Your value	Minimum directory size	More information
TCP/IP port address for communications with the server	1500		Not applicable.	Ensure that this port is available when you install and configure the operating system. The port number can be a number in the range 1024 - 32767.

Item	Default value	Your value	Minimum directory size	More information
Directory for the server instance	<p>AIX Linux /home/tsminst1/tsminst1</p> <p>Windows C:\tsminst1</p>		<p>AIX 50 GB.</p> <p>Linux Windows 25 GB.</p>	If you change the value for the server instance directory from the default, also modify the DB2® instance owner value in Table 2.
Directory for server installation	<ul style="list-style-type: none"> • AIX Linux / • Windows C: 		<p>AIX Available space that is required for the directory: 5 GB.</p> <p>Linux Windows Minimum space that is required for the directory: 30 GB</p>	
Directory for server installation	/usr		AIX Available space that is required for the directory: 5 GB.	
Directory for server installation	AIX /var		AIX Available space that is required for the directory: 5 GB.	
Directory for server installation	AIX /tmp		AIX Available space that is required for the directory: 5 GB.	
Directory for server installation	AIX /opt		AIX Available space that is required for the directory: 10 GB.	
Directory for the active log	<p>AIX Linux /tsminst1/TSMalog</p> <p>Windows C:\tsminst1\TSMalog</p>		128 GB.	When you create the active log during the initial configuration of the server, set the size to 128 GB.
Directory for the archive log	<p>AIX Linux /tsminst1/TSMarchlog</p> <p>Windows C:\tsminst1\TSMarchlog</p>		3 TB.	
Directories for the database	<p>AIX Linux /tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03</p> <p>Windows C:\tsminst1\TSMdbspace00 C:\tsminst1\TSMdbspace01 C:\tsminst1\TSMdbspace02 C:\tsminst1\TSMdbspace03</p>		For instructions about calculating space requirements, see Hardware requirements.	Create four file systems for the database.

Item	Default value	Your value	Minimum directory size	More information
Directories for storage	<p>AIX Linux</p> /tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ... <p>Windows C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...</p>		Determine the minimum total capacity for all directories by using the following calculation: Daily percentage of ingested data that is written to disk + 20% = Minimum total capacity	The preferred method is to define at least one directory for each tape device.

Table 2. Worksheet for IBM Spectrum Protect configuration

Item	Default value	Your value	More information
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 1 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	<p>AIX Linux passwOrd</p> <p>Windows pAssWOrd</p>		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	<p>AIX Linux tsmsrvrs</p>		
Server name	The default value for the server name is the system host name.		
Server password	passwOrd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
Administrator ID password	passwOrd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Item	Default value	Your value	More information
Schedule start time	23:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p> <p>In this guide, the suggested time to start client backup operations is 23:00.</p>

Table 3. Worksheet for tape configuration

Item	Default value	Your value	More information
Robotic device files	<p>IBM® devices with an IBM tape device driver:</p> <ul style="list-style-type: none"> • AIX /dev/smcX • Linux /dev/IBMchangerX • Windows ChangerX <p>Non-IBM devices with an IBM Spectrum Protect device driver:</p> <ul style="list-style-type: none"> • AIX /dev/lbX • Linux /dev/tsm SCSI/lbX • Windows lbA.B.C.D 		<p>To manually define the library device files, use the following commands:</p> <ul style="list-style-type: none"> • DEFINE LIBRARY • DEFINE DRIVE • DEFINE PATH <p>For SCSI, you can use the PERFORM LIBACTION command to define all drives and their paths for a single library in one step. To use this command to define all drives and paths, the SANDISCOVERY option must be supported and enabled.</p>
Tape drives	<p>IBM devices with an IBM tape device driver:</p> <ul style="list-style-type: none"> • AIX /dev/rmtX • Linux /dev/IBMtapeX • Windows TapeX <p>Non-IBM devices with an IBM Spectrum Protect device driver:</p> <ul style="list-style-type: none"> • AIX /dev/mtX • Linux /dev/tsm SCSI/mtX • Windows mtA.B.C.D 		

Planning for disk storage

Choose the most effective storage technology for IBM Spectrum Protect™ components to ensure efficient server performance and operations.

Storage hardware devices have different capacity and performance characteristics, which determine how they can be used effectively with IBM Spectrum Protect. For general guidance about selecting the appropriate storage hardware and setup for your solution, review the following guidelines.

Database, active log, and archive log

- Use a solid-state disk (SSD) or a fast, 15,000 rpm disk for the IBM Spectrum Protect database and active log.
- When you create arrays for the database, use RAID level 5.
- Use separate disks for archive log and database backup storage.

Storage pool

Use RAID level 6 for storage pool arrays to add protection against double drive failures when you use large disk types.

- Planning the storage arrays
Prepare for disk storage configuration by planning for RAID arrays and volumes, according to the size of your IBM Spectrum Protect system.

Planning for tape storage




Determine which tape devices to use and how to configure them. To optimize system performance, plan to use fast, high-capacity tape devices. Provision enough tape drives to meet your business requirements.

- Supported tape devices and libraries
The server can use a wide range of tape devices and libraries. Select tape devices and libraries that meet your business requirements.
- Supported tape device configurations
Review the information about local area networks (LAN) and storage area networks (SAN). To optimize data movement, plan to configure LAN-free data movement. In addition, consider whether to use library sharing.
- Required definitions for tape storage devices
Before the IBM Spectrum Protect™ server can use a tape device, you must configure the device to the operating system and to the server. As part of the planning process, determine which definitions are required for your tape storage devices.
- Planning the storage pool hierarchy
Plan the storage pool hierarchy to ensure that data is migrated daily from disk to tape. The migration releases space on the disk device and moves the data to tape for long-term retention. In this way, you can take advantage of the scalability, cost efficiency, and security features of tape storage.
- Offsite data storage
To facilitate data recovery and as part of your disaster recovery strategy, store tape copies offsite.

Supported tape devices and libraries

The server can use a wide range of tape devices and libraries. Select tape devices and libraries that meet your business requirements.

For a list of supported devices and valid device class formats, see the website for your operating system:

-   Supported devices for AIX and Windows
-  Supported devices for Linux

For more information about storage devices and storage objects, see Types of storage devices.

Each device that is defined to IBM Spectrum Protect™ is associated with one *device class*. The device class specifies the device type and media management information, such as recording format, estimated capacity, and labeling prefixes.

A *device type* identifies a device as a member of a group of devices that share similar media characteristics. For example, the LTO device type applies to all generations of LTO tape drives.

A device class for a tape drive must also specify a library. A *physical library* is a collection of one or more drives that share similar media-mounting requirements. That is, the drive can be mounted by an operator or by an automated mounting mechanism.

A *library object definition* specifies the library type and other characteristics that are associated with that library type.

The following table lists the preferred library types for an IBM Spectrum Protect Version 8.1.5 tape solution.

Table 1. Library types for an IBM Spectrum Protect 8.1.5 tape solution

Library type	Description	More information
SCSI	<p>A SCSI library is controlled through a SCSI interface, attached either directly to the server's host by using SCSI cabling or by a storage area network. A robot or other mechanism automatically handles tape volume mounts and dismounts.</p> <p>If you create different drive types for a SCSI library, you create multiple logical libraries that cannot be split between different types of drives. A SCSI library can contain drives of mixed technologies, including LTO Ultrium and digital linear tape (DLT) drives. For example:</p> <ul style="list-style-type: none"> • The Oracle StorageTek L700 library • The IBM® 3592 tape device 	<p>Configuring libraries for use by a server</p> <p>Restrictions apply when you mix different generations of media and drives. For more information, see:</p> <ul style="list-style-type: none"> • Mixing generations of 3592 drives and media in a single library • Mixing LTO drives and media in a library
Shared	<p>Shared libraries are logical libraries that are represented by SCSI. The library is controlled by the IBM Spectrum Protect server that is configured as a library manager.</p> <p>IBM Spectrum Protect servers that use the SHARED library type are library clients to the library manager server. Shared libraries reference a library manager.</p>	

Supported tape device configurations

Review the information about local area networks (LAN) and storage area networks (SAN). To optimize data movement, plan to configure LAN-free data movement. In addition, consider whether to use library sharing.

Select the device configuration that meets your business requirements.

- LAN-based and LAN-free data movement
You can move data between clients and storage devices that are attached to a local area network (LAN), or to storage devices that are attached to a storage area network (SAN), known as LAN-free data movement.
- Library sharing
You can optimize the efficiency of your tape solution by configuring library sharing. Library sharing allows multiple IBM Spectrum Protect™ servers to use the same tape library and drives on a storage area network (SAN) and to improve backup and recovery performance and tape hardware utilization.
- LAN-free data movement
IBM Spectrum Protect provides the capability for a client, through a storage agent, to directly back up and restore data to a tape library on a SAN. This type of data movement is also known as LAN-free data movement.
- Mixed device types in libraries
IBM Spectrum Protect supports mixing different device types within a single automated library, if the library can distinguish among the different media for the different device types. To simplify the configuration process, do not plan to mix different device types within a library. If you must mix device types, review the restrictions.

LAN-based and LAN-free data movement

You can move data between clients and storage devices that are attached to a local area network (LAN), or to storage devices that are attached to a storage area network (SAN), known as LAN-free data movement.

In a conventional LAN configuration, one or more tape libraries are associated with a single IBM Spectrum Protect™ server. LAN-free data movement makes LAN bandwidth available for other uses and decreases the load on the IBM Spectrum Protect server.

In a LAN configuration, client data, email, terminal connection, application program, and device control information must be handled by the same network. Device control information and client backup and restore data flow across the LAN.

A SAN is a dedicated storage network that can improve system performance.

By using IBM Spectrum Protect in a SAN, you benefit from the following functions:

- Sharing storage devices among multiple IBM Spectrum Protect servers.
Restriction: A storage device with the GENERICTAPE device type cannot be shared among servers.
- Moving IBM Spectrum Protect client data directly to storage devices (LAN-free data movement) by configuring a storage agent on the client system.

In a SAN, you can share tape drives and libraries that are supported by the IBM Spectrum Protect server, including most SCSI tape devices.

When IBM Spectrum Protect servers share a SCSI tape, one server, the *library manager*, owns and controls the device. The storage agents, along with other IBM Spectrum Protect servers that share this library are *library clients*. A library client requests shared library resources, such as drives or media, from the library manager, but uses the resources independently. The library manager coordinates the access to these resources. IBM Spectrum Protect servers that are defined as library clients use server-to-server communications to contact the library manager and request device service. Data moves over the SAN between each server and the storage device.

Requirement: If you define a library manager server that is shared with the IBM Spectrum Protect server, the SANDISCOVERY option must be set to ON. By default, this option is set to OFF.

IBM Spectrum Protect servers use the following features when sharing an automated library:

Partitioning of the volume inventory

The inventory of media volumes in the shared library is partitioned among servers. Either one server owns a particular volume, or the volume is in the global scratch pool. No server owns the scratch pool.

Serialized drive access

Only one server accesses each tape drive at a time. Drive access is serialized. IBM Spectrum Protect controls drive access so that servers do not dismount other servers' volumes or write to drives where other servers mount their volumes.

Serialized mount access

The library autochanger completes a single mount or dismount operation at a time. The library manager completes all mount operations to provide this serialization.

Library sharing

You can optimize the efficiency of your tape solution by configuring library sharing. Library sharing allows multiple IBM Spectrum Protect™ servers to use the same tape library and drives on a storage area network (SAN) and to improve backup and recovery performance and tape hardware utilization.

When IBM Spectrum Protect servers share a library, one server is set up as the library manager and controls library operations such as mount and dismount. The library manager also controls volume ownership and the library inventory. Other servers are set up as library clients and use server-to-server communications to contact the library manager and request resources.

Library clients must be at the same or an earlier version than the library manager server. A library manager cannot support library clients that are at a later version. For more information, see Storage-agent and library-client compatibility with an IBM Spectrum Protect server.


LAN-free data movement

IBM Spectrum Protect™ provides the capability for a client, through a storage agent, to directly back up and restore data to a tape library on a SAN. This type of data movement is also known as LAN-free data movement.

Restriction: Centera storage devices cannot be targets for LAN-free operations.

Figure 1 shows a SAN configuration in which a client directly accesses a tape to read or write data.

Figure 1. LAN-free data movement



LAN-free data movement requires the installation of a storage agent on the client system. The server maintains the database and recovery log, and acts as the library manager to control device operations. The storage agent on the client handles the data transfer to the device on the SAN. This implementation frees up bandwidth on the LAN that would otherwise be used for client data movement.

Mixed device types in libraries

IBM Spectrum Protect™ supports mixing different device types within a single automated library, if the library can distinguish among the different media for the different device types. To simplify the configuration process, do not plan to mix different device types within a library. If you must mix device types, review the restrictions.

Libraries with this capability are models that have built-in mixed drives, or that support the addition of mixed drives. For information about specific models, see the manufacturer's documentation. To learn about libraries that were tested on IBM Spectrum Protect with mixed device types, see the information for your operating system:

- IBM Spectrum Protect Supported Devices for AIX®, HP-UX, Solaris, and Windows
- IBM Spectrum Protect Supported Devices for Linux

For example, you can have LTO Ultrium drives and IBM TS4500 drives in a single library that is defined to the IBM Spectrum Protect server.

- Different media generations in a library
The IBM Spectrum Protect server allows mixed device types in an automated library, but the mixing of different generations of the same type of drive is generally not supported. New drives cannot write to the older media formats, and old drives cannot read new formats. LTO Ultrium drives are an exception to this rule.
- Mixed media and storage pools
You can optimize the efficiency of your tape solution by not mixing media formats in a storage pool. Instead of mixing

formats, map each unique media format to a separate storage pool by using its own device class. This restriction also applies to LTO formats.

Different media generations in a library

The IBM Spectrum Protect™ server allows mixed device types in an automated library, but the mixing of different generations of the same type of drive is generally not supported. New drives cannot write to the older media formats, and old drives cannot read new formats. LTO Ultrium drives are an exception to this rule.

If the new drive technology cannot write to media that is formatted by older generation drives, the older media must be marked read-only to avoid problems for server operations. Also, the older drives must be removed from the library, or the definitions of the older drives must be removed from the server. For example, the IBM Spectrum Protect server does not support the use of Oracle StorageTek 9940A drives with 9940B drives in combination with other device types in a single library.

In general, IBM Spectrum Protect does not support mixing generations of LTO Ultrium drives and media. However, the following mixtures are supported:

- LTO Ultrium Generation 3 (LTO-3) with LTO Ultrium Generation 4 (LTO-4)
- LTO Ultrium Generation 4 (LTO-4) with LTO Ultrium Generation 5 (LTO-5)
- LTO Ultrium Generation 5 (LTO-5) with LTO Ultrium Generation 6 (LTO-6)
- LTO Ultrium Generation 6 (LTO-6) with LTO Ultrium Generation 7 (LTO-7)
- LTO Ultrium Generation 7 (LTO-7) media with LTO Ultrium Generation 8 (LTO-8 and LTO-M8) media in a library with LTO-8 tape drives or a library with mixed LTO-8 and LTO-7 tape drives

The server supports these mixtures because the different drives can read and write to the different media. If you plan to upgrade all drives to Generation 4 (or Generation 5, 6, 7, or 8), you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4 (or Generation 5, 6, 7, or 8) drives and paths.

Restrictions that apply to mixing LTO Ultrium tape drives and media

- LTO-5 drives can read only LTO-3 media. If you are mixing LTO-3 with LTO-5 drives and media in a single library, you must mark the LTO-3 media as read-only. You must check out all LTO-3 scratch volumes.
- LTO-6 drives can read only LTO-4 media. If you are mixing LTO-4 with LTO-6 drives and media in a single library, you must mark the LTO-4 media as read-only. You must check out all LTO-4 scratch volumes.
- LTO-7 drives can read only LTO-5 media. If you are mixing LTO-5 with LTO-7 drives and media in a single library, you must mark the LTO-5 media as read-only. You must check out all LTO-5 scratch volumes.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 and LTO-8 drives and media in a single library, you must partition the library into two libraries. One library has only LTO-8 drives and media and the other has LTO-6 drives and media.

Restrictions that apply to mixed generation LTO Ultrium tape drives in a library

You must use tape cartridges that are an earlier generation than the tape drive. A later generation tape drive can read and write data to an earlier generation tape cartridge. For an example, if a library has LTO-7 and LTO-6 tape drives, you must use LTO-6 tape cartridges. Both the LTO-7 and LTO-6 tape drives can read and write data to LTO-6 tape cartridges.

Restrictions that apply to mixed generation LTO Ultrium tape cartridges in a library

You must use a tape cartridge that is the same generation as the tape drive, or one generation earlier. For example, if a library has LTO-7 tape drives, you can use LTO-7 tape cartridges or mixed LTO-7 and LTO-6 tape cartridges. If this library has LTO-7, LTO-6, and LTO-5 tape cartridges, you must change the access mode to READONLY for the LTO-5 tape cartridges.

To learn about additional considerations when you mix LTO Ultrium generations, see Defining LTO device classes.

When you use IBM Spectrum Protect, you cannot mix drives that are 3592, TS1130, TS1140, TS1150, and later drive generations. Use one of three special configurations. For details, see Defining 3592 device classes.

If you plan to encrypt volumes in a library, do not mix media generations in the library.

Mixed media and storage pools

You can optimize the efficiency of your tape solution by not mixing media formats in a storage pool. Instead of mixing formats, map each unique media format to a separate storage pool by using its own device class. This restriction also applies to LTO formats.

Multiple storage pools and their device classes of different types can point to the same library that can support them as described in Different media generations in a library.

You can migrate to a new generation of a media type within the same storage pool by following these steps:

1. Replace all older drives with the newer generation drives within the library. The drives should be mixed.
2. Mark the existing volumes with the older formats read-only if the new drive cannot append those tapes in the old format. If the new drive can write to the existing media in their old format, this is not necessary, but Step 1 is still required. If it is necessary to keep different drive generations that are read but not write compatible within the same library, use separate storage pools for each.

Required definitions for tape storage devices

Before the IBM Spectrum Protect™ server can use a tape device, you must configure the device to the operating system and to the server. As part of the planning process, determine which definitions are required for your tape storage devices.

Tip: You can use the PERFORM LIBACTION command to simplify the process when you add devices to SCSI and VTL library types.

Table 1 summarizes the definitions that are required for different device types.

Table 1. Required definitions for storage devices

Device	Device types	Required definitions			
		Library	Drive	Path	Device class
Magnetic disk	DISK	—	—	—	Yes ¹
	FILE ²	—	—	—	Yes
	AIX Windows CENTERA Linux CENTERA ³	—	—	—	Yes
Tape	<ul style="list-style-type: none"> • 3590 • 3592 • DLT • LTO • NAS • VOLSAFE AIX Windows GENERICTAPE ECARTRIDGE ⁴	Yes	Yes	Yes	Yes
Removable media (file system)	REMOVABLEFILE	Yes	Yes	Yes	Yes

1. The DISK device class exists at installation and cannot be changed.
2. FILE libraries, drives, and paths are required for sharing with storage agents.
3. Linux The CENTERA device type is available only for Linux x86_64 systems.
4. The ECARTRIDGE device type is for Oracle StorageTek cartridge tape drives such as 9840 and T10000 drives.

Planning the storage pool hierarchy

Plan the storage pool hierarchy to ensure that data is migrated daily from disk to tape. The migration releases space on the disk device and moves the data to tape for long-term retention. In this way, you can take advantage of the scalability, cost efficiency, and security features of tape storage.

Before you begin

The storage pool hierarchy helps to manage the flow of data. To understand the data flow, review Figure 1.
Figure 1. Tape solution



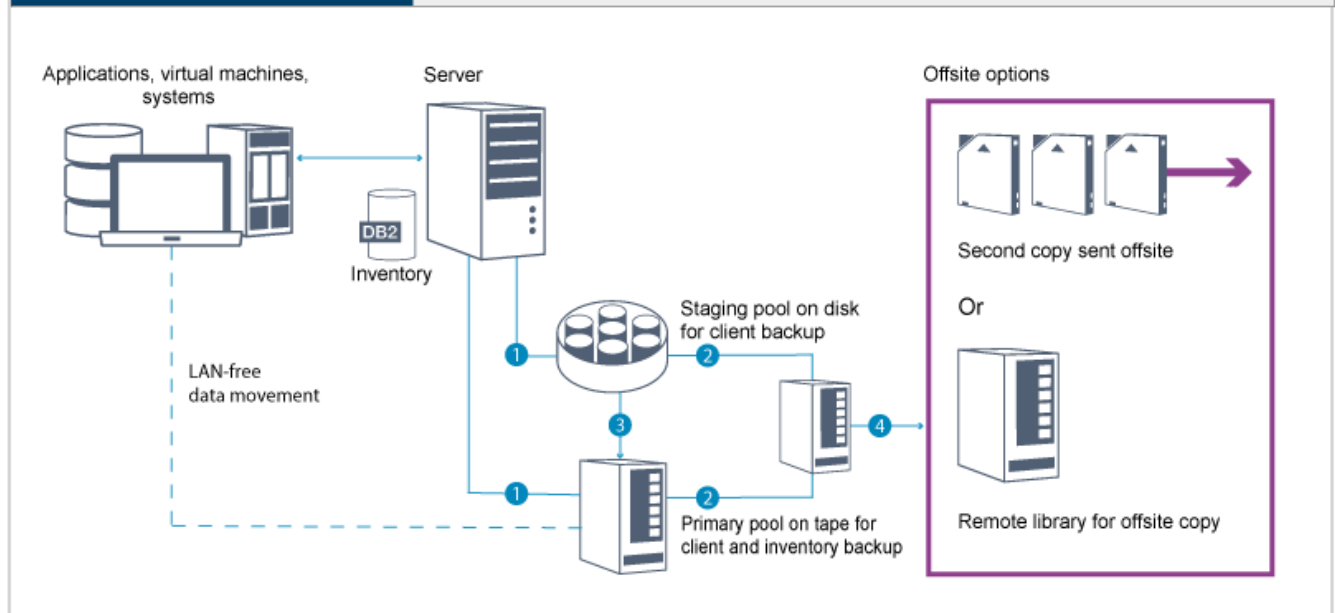
Tape

✓ Ideal for long-term retention

✓ Low-cost scalability

✓ Disk staging to primary tape pool

✓ Optimized for SAN



The following steps correspond to the numbers in the figure:

1. The server receives data from clients (applications, virtual machines, or systems) and stores the data on primary storage pools. Depending on the client type, the data is stored on a primary storage pool on disk or tape.
2. The data on disk and tape is backed up to a copy storage pool on tape.
3. Data in the primary storage pool on disk is migrated daily to the primary storage pool on tape.
4. Data from the copy storage pool on tape is moved offsite to support long-term retention and disaster recovery.

Procedure

To plan the storage pool hierarchy, answer the following questions:

- a. Which clients should back up data to disk, and which clients should back up data to tape?
 - The preferred method is to back up clients that host large objects, such as databases, to tape.
 - The preferred method is to back up all other clients to disk.
 - Virtual machine (VM) clients can be backed up to disk or tape. The preferred method is to back up a VM client to a separate disk storage pool, which is not migrated to tape. If you must migrate a VM client to tape, create a smaller disk storage pool to hold the VMware control files. This smaller disk storage pool cannot be allowed to migrate to tape. For more information about backing up a VM client to tape, see Tape media guidelines and technote 1239546.

Tip: If many clients must back up data to a single storage pool, consider using a storage pool on disk because you can specify many mount points. You can specify a maximum value of 999 for the MAXNUMMP parameter on the REGISTER NODE command.

- b. What are the considerations for specifying the capacity of disk-based storage pools?

At minimum, plan enough capacity to store data from a single day of backup operations. The preferred method is to plan enough capacity to store data from two days' worth of backup operations and add a 20% buffer.

- c. What are the considerations for specifying the device class for the disk-based storage pool?

The preferred method is to specify a FILE device class. Set the MOUNTLIMIT parameter to 4000. Also, ensure that the node has a sufficiently high number of mount points, which you can specify by using the MAXNUMMP parameter on the REGISTER NODE command.

- d. Should data deduplication be specified for the disk storage pool?

No, because the data is stored on disk for only one day before the data is migrated to tape.

- e. Should automatic migration of data be specified based on a migration threshold?

No. Instead, plan to schedule daily migration by using the MIGRATE STGPOOL command. (To prevent automatic migration based on the migration threshold, specify a value of 100 for the HIGHMIG parameter and 0 for the LOWMIG parameter when you issue the DEFINE STGPOOL command.)

f. Should a migration delay be specified?

The preferred method is to specify migration from disk to tape daily, and not specify a migration delay, which requires additional planning. For more information about migration delays, see *Migrating files in a storage pool hierarchy*.

g. How can the number of tape drives be calculated?

- i. Determine the native data transfer rate of the drive by reviewing the manufacturer's documentation. To obtain an estimate of the sustained data transfer rate in your storage environment, subtract 30% from the native data transfer rate.
- ii. Calculate the required rate of data ingestion by the server. Then, divide that figure by the sustained data transfer rate of a single tape device. The result is the minimum number of drives to support data ingestion.
- iii. Calculate the number of mount points that are required by clients that back up data to tape, including those clients that use multiple sessions. You can distribute the mount points over the backup window, keeping in mind that clients are likely backing up large objects, which might use most of the window.
- iv. Calculate the performance requirements *and* mount points that are required for maintenance tasks, such as disk-to-tape migration and tape-to-tape copies. By backing up data to tape, you can avoid migration processing, but making tape-to-tape copies will double the tape drive requirement.
- v. Calculate the number of additional drives that might be required, for example:
 - If a tape drive malfunctions, the issue impacts the number of available mount points and the ingestion rate. Consider provisioning spare drives. For example, if you require five tape drives for normal operations, consider provisioning two spare drives.
 - Restore and retrieve operations might require additional tape drives if you plan to run the operations simultaneously with data ingestion and maintenance operations. If necessary, provision additional tape drives and ensure that they are unused when you start the restore or retrieve operations.

h. What alternatives are available for optimizing restore operations?

You can use collocation to improve system performance and optimize data organization. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored:

- For disk-based storage pools, the preferred method is to use collocation by node. The server stores the data for the node on as few volumes as possible.
- For tape-based storage pools, the preferred method is to use collocation by group. Collocation by group results in a reduction of unused tape capacity, which allows for more collocated data on individual tapes.

For more information about collocation, see *Optimizing operations by enabling collocation of client files*.

If you are an experienced system administrator, you might plan additional actions to optimize restore operations. See *Optimizing restore operations for clients*, *File backup techniques*, and *MOVE NODEDATA* (Move data by node in a sequential access storage pool).

Offsite data storage

To facilitate data recovery and as part of your disaster recovery strategy, store tape copies offsite.

Use the disaster recovery manager (DRM) function to configure and automatically generate a disaster recovery plan that contains the information, scripts, and procedures that are required to automatically restore the server and recover client data after a disaster. Choose from one of the following offsite data storage options as a disaster recovery strategy to protect tape copies:

Offsite vaulting from a single production site

Storage volumes, such as tape cartridges and media volumes, are vaulted at an offsite location. A courier transports the data from the offsite storage facility to the recovery site. If a disaster occurs, the volumes are sent back to the production site after hardware and the IBM Spectrum Protect™ server are restored.

Offsite vaulting with a recovery site

A courier moves storage volumes from the production site to an offsite storage facility. By having a dedicated recovery site, you can reduce recovery time compared to the single production site. However, this option increases the cost of disaster recovery because more hardware and software must be maintained. For example, the recovery site must have compatible tape devices and IBM Spectrum Protect server software. Before the production site can be recovered, the hardware and software at the recovery site must be set up and running.

Electronic vaulting

To use electronic vaulting as a disaster recovery strategy, the recovery site must have a running IBM Spectrum Protect server. Critical data is vaulted electronically from the production site to the recovery site. DRM is also used for offsite

vaulting of noncritical data. Electronic vaulting moves critical data offsite faster and more frequently than traditional courier methods. Recovery time is reduced because critical data is already stored at the recovery site. However, because the recovery site runs continuously, the cost of the disaster recovery strategy is more expensive than offsite vaulting.

Related concepts:

Preparing for and recovering from a disaster by using DRM

Planning for security

Plan to protect the security of systems in the IBM Spectrum Protect™ solution with access and authentication controls, and consider encrypting data and password transmission.

- **Planning for administrator roles**
Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect solution.
- **Planning for secure communications**
Plan for protecting communications among the IBM Spectrum Protect solution components.
- **Planning for storage of encrypted data**
Determine whether your company requires stored data to be encrypted, and choose the method that best suits your needs.
- **Planning firewall access**
Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect solution to work.

Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the IBM Spectrum Protect™ solution.

You can assign one of the following levels of authority to administrators:

System

Administrators with system authority have the highest level of authority. Administrators with this level of authority can complete any task. They can manage all policy domains and storage pools, and grant authority to other administrators.

Policy

Administrators who have policy authority can manage all of the tasks that are related to policy management. This privilege can be unrestricted, or can be restricted to specific policy domains.

Storage

Administrators who have storage authority can allocate and control storage resources for the server.

Operator

Administrators who have operator authority can control the immediate operation of the server and the availability of storage media such as tape libraries and drives.

The scenarios in Table 1 provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

Table 1. Scenarios for administrator roles

Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	<ul style="list-style-type: none"> • System authority: 1 administrator ID
An administrator for multiple servers also manages the overall system. Several other administrators manage their own storage pools.	<ul style="list-style-type: none"> • System authority on all servers: 1 administrator ID for the overall system administrator • Storage authority for designated storage pools: 1 administrator ID for each of the other administrators
An administrator manages 2 servers. Another person helps with the administration tasks. Two assistants are responsible for helping to ensure that important systems are backed up. Each assistant is responsible for monitoring the scheduled backups on one of the IBM Spectrum Protect servers.	<ul style="list-style-type: none"> • System authority on both servers: 2 administrator IDs • Operator authority: 2 administrator IDs for the assistants with access to the server that each person is responsible for

Related tasks:

Managing administrators

Planning for secure communications

Plan for protecting communications among the IBM Spectrum Protect™ solution components.

Determine the level of protection that is required for your data, based on regulations and business requirements under which your company operates.

If your business requires a high level of security for passwords and data transmission, plan on implementing secure communication with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocols.

TLS and SSL provide secure communications between the server and client, but can affect system performance. To improve system performance, use TLS for authentication without encrypting object data. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option for client-to-server communication, and the UPDATE SERVER=SSL parameter for server-to-server communication. Beginning in V8.1.2, TLS is used for authentication by default. If you decide to use TLS to encrypt entire sessions, use the protocol only for sessions where it is necessary and add processor resources on the server to manage the increase in network traffic. You can also try other options. For example, some networking devices such as routers and switches provide the TLS or SSL function.

You can use TLS and SSL to protect some or all of the different possible communication paths, for example:

- Operations Center: browser to hub; hub to spoke
- Client to server
- Server to server: node replication

Related tasks:

Configuring secure communications with Transport Layer Security

Planning for storage of encrypted data

Determine whether your company requires stored data to be encrypted, and choose the method that best suits your needs.

Table 1. Selecting a data encryption method

Business requirement	Encryption method	Additional information
Protect data at the client level.	IBM Spectrum Protect™ client encryption	You can encrypt data at the file level by using an include/exclude list. In this way, you can maintain a high degree of control over which data is encrypted. Extra computing resources are required at the client that might affect the performance of backup and restore processes. For more information about this method, see IBM Spectrum Protect client encryption.
Protect data in storage pool volumes on a tape drive.	Application method	When you use the Application method, IBM Spectrum Protect manages the encryption keys to protect data in storage pool volumes. You must take extra care to secure database backups because the encryption keys are stored in the server database. Without access to database backups and matching encryption keys, you cannot restore your data. You cannot use this method to encrypt database backups, exported data, or backup sets. For more information about the Application method, see Tape encryption methods.
Protect data on a tape drive.	Library method	When you use the Library method, the library manages encryption keys. You can encrypt both data in storage pools and other data on a tape drive. You can control which volumes are encrypted by using their bar code serial numbers. For more information about the Library method, see Tape encryption methods.
Protect data on a tape drive.	System method	When you use the System method, a device driver or the AIX operating system manages encryption. This encryption method is available only on the AIX® operating system. You can encrypt both data in storage pools and other data on a tape drive. For more information about the System method, see Tape encryption methods.

Planning firewall access

Determine the firewalls that are set and the ports that must be open for the IBM Spectrum Protect™ solution to work.

Table 1 describes the ports that are used by the server, client, and Operations Center.

Table 1. Ports that are used by the server, client, and Operations Center

Item	Default	Direction	Description
Base port (TCPPOINT)	1500	Outbound/inbound	Each server instance requires a unique port. You can specify an alternative port number. The TCPPOINT option listens for both TCP/IP and SSL-enabled sessions from the client. You can use the TCPADMINPORT option and ADMINONCLIENTPORT option to set port values for administrative client traffic.
SSL-only port (SSLTCPPOINT)	No default	Outbound/inbound	This port is used if you want to restrict communication on the port to SSL-enabled sessions only. A server can support both SSL and non-SSL communication by using the TCPPOINT or TCPADMINPORT options.
SMB	45	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SSH	22	Inbound/outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SMTP	25	Outbound	This port is used to send email alerts from the server.
Replication	No default	Outbound/inbound	The port and protocol for the outbound port for replication are set by the DEFINE SERVER command that is used to set up replication. The inbound ports for replication are the TCP ports and SSL ports are specified for the source server on the DEFINE SERVER command.
Client schedule port	Client port: 1501	Outbound	The client listens on the port that is named and communicates the port number to the server. The server contacts the client if server prompted scheduling is used. You can specify an alternative port number in the client options file.
Long-running sessions	KEEPALIVE setting: YES	Outbound	When the KEEPALIVE option is enabled, keepalive packets are sent during client/server sessions to prevent the firewall software from closing long-running, inactive connections.
Operations Center	HTTPS: 11090	Inbound	These ports are used for the Operations Center web browser. You can specify an alternative port number.
Client management service port	Client port: 9028	Inbound	If you plan to use IBM Spectrum Protect client management services, the client management service port must be accessible from the Operations Center. Ensure that firewalls cannot prevent connections. The client management service uses the TCP port of the server for the client node for authentication by using an administrative session.

Related tasks:

- 🔗 [Collecting diagnostic information with IBM Spectrum Protect client management services](#)

Related reference:

- 🔗 [ADMINONCLIENTPORT server option](#)
- 🔗 [DEFINE SERVER \(Define a server for server-to-server communications\)](#)
- 🔗 [TCPADMINPORT server option](#)
- 🔗 [TCPPOINT server option](#)

Implementation of a tape-based data protection solution

Implement the tape-based solution, which uses disk-to-disk-to-tape backup and disk staging to optimize storage. By implementing the tape solution, you can enable long-term data retention and achieve low-cost scalability.




Tape

✓ Ideal for long-term retention

✓ Low-cost scalability

✓ Disk staging to primary tape pool

✓ Optimized for SAN



Tip: The described solution does not include node replication. However, if you want to use node replication to back up a storage pool from disk to disk, ensure that the replication operation is completed before data is migrated from disk to tape. You can also use node replication to back up a storage pool on a local tape device to a copy storage pool on a local tape device.

Implementation roadmap

The following steps are required to set up a tape-based solution.

1. Set up the system.
2. Install the server and the Operations Center.
3. Configure the server and the Operations Center.
4. Attach tape devices for the server.
5. Configure tape libraries for use by the server.
6. Set up a storage pool hierarchy.
7. Install and configure clients.
8. Configure LAN-free data movement.
9. Select an encryption method and configure encryption.
10. Set up tape storage operations.
11. Complete the implementation.

Setting up the system

To set up the system, you must first configure your disk storage hardware and the server system for IBM Spectrum Protect™.

About this task

Tip: Procedures for setting up the server and the disk storage system are described. To get started with setting up tape devices, see [Attaching tape devices for the server](#).

- **Configuring the storage hardware**
To optimize disk storage, review the guidelines for setting up disk storage with IBM Spectrum Protect. Then, provide a connection between the server and the disk storage devices and complete other configuration tasks.
- **Installing the server operating system**
Install the operating system on the server system and ensure that IBM Spectrum Protect server requirements are met. Adjust operating system settings as directed.
- **Configuring multipath I/O**
You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for

detailed instructions.

- Creating the user ID for the server
Create the user ID that owns the IBM Spectrum Protect server instance. You specify this user ID when you create the server instance during initial configuration of the server.
- Preparing file systems for the server
You must complete file system configuration for the disk storage to be used by the server.

Configuring the storage hardware

To optimize disk storage, review the guidelines for setting up disk storage with IBM Spectrum Protect. Then, provide a connection between the server and the disk storage devices and complete other configuration tasks.

Before you begin

For guidelines about setting up disk storage, see Checklist for storage pools on DISK or FILE

Procedure

1. Provide a connection between the server and the storage devices by following these guidelines:
 - Use a switch or direct connection for Fibre Channel connections.
 - Consider the number of ports that are connected and account for the amount of bandwidth that is needed.
 - Consider the number of ports on the server and the number of host ports on the disk system that are connected.
2. Verify that device drivers and firmware for the server system, adapters, and operating system are current and at the recommended levels.
3. Configure storage arrays. Make sure that you planned properly to ensure optimal performance. For more information, see Planning for disk storage.
4. Ensure that the server system has access to disk volumes that are created. Complete the following steps:
 - a. If the system is connected to a Fibre Channel switch, zone the server to see the disks.
 - b. Map all of the volumes to tell the disk system that this specific server is allowed to see each disk.
5. Ensure that tape and disk devices use different Host Bus Adapter (HBA) ports. Control tape and disk I/O by using the SAN.

Related tasks:

Configuring multipath I/O

Installing the server operating system

Install the operating system on the server system and ensure that IBM Spectrum Protect™ server requirements are met. Adjust operating system settings as directed.

- Installing on AIX systems
Complete the following steps to install AIX® on the server system.
- Installing on Linux systems
Complete the following steps to install Linux x86_64 on the server system.
- Installing on Windows systems
Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect server.

Installing on AIX systems

Complete the following steps to install AIX® on the server system.

Procedure

1. Install AIX Version 7.1, TL4, SP2, or later according to the manufacturer instructions.
2. Configure your TCP/IP settings according to the operating system installation instructions.
3. Open the /etc/hosts file and complete the following actions:
 - Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7 server.yourdomain.com server
```

- o Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1 localhost
```

4. Enable AIX I/O completion ports by issuing the following command:

```
chdev -l iocp0 -P
```

Server performance can be affected by the Olson time zone definition.

5. To optimize performance, change your system time zone format from Olson to POSIX. Use the following command format to update the time zone setting:

```
chtz=local_timezone,date/time,date/time
```

For example, if you lived in Tucson, Arizona, where Mountain Standard Time is used, you would issue the following command to change to the POSIX format:

```
chtz MST7MDT,M3.2.0/2:00:00,M11.1.0/2:00:00
```

6. Add an entry in the .profile of the instance user so that the following environment is set:

```
export MALLOCOPTIONS=multiheap:16
```

Tip: If the instance user is not available, complete this step later, when the instance user becomes available.

7. Set the system to create full application core files. Issue the following command:

```
chdev -l sys0 -a fullcore=true -P
```

8. For communications with the server and Operations Center, make sure that the following ports are open on any firewalls that might exist:

- o For communications with the server, open port 1500.
- o For secure communications with the Operations Center, open port 11090 on the hub server.

If you are not using the default port values, make sure that the ports that you are using are open.

9. Enable TCP high-performance enhancements. Issue the following command:

```
no -p -o rfc1323=1
```

10. For optimal throughput and reliability, bond four 10 Gb Ethernet ports together. Use the System Management Interface Tool (SMIT) to bond the ports together by using Etherchannel. The following settings were used during testing:

```
mode          8023ad
auto_recovery yes          Enable automatic recovery after failover
backup_adapter NONE       Adapter used when whole channel fails
hash_mode     src_dst_port Determines how outgoing adapter is chosen
interval      long        Determines interval value for IEEE
802.3ad mode
mode          8023ad      EtherChannel mode of operation
netaddr       0           Address to ping
no_loss_failover yes     Enable lossless failover after ping
failure
num_retries   3           Times to retry ping before failing
retry_time    1           Wait time (in seconds) between pings
use_alt_addr  no          Enable Alternate EtherChannel Address
use_jumbo_frame no       Enable Gigabit Ethernet Jumbo Frames
```

11. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 1. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 1. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf

User limit type	Setting	Value	Command to query value
Maximum number of open files	nofile	65536	<code>ulimit -Hn</code>
Maximum amount of processor time in seconds	cpu	Unlimited	<code>ulimit -Ht</code>
Maximum number of user processes	nproc	16384	<code>ulimit -Hu</code>

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Linux systems

Complete the following steps to install Linux x86_64 on the server system.

Before you begin

The operating system will be installed on the internal hard disks. Configure the internal hard disks by using a hardware RAID 1 array. For example, if you are configuring a small system, the two 300 GB internal disks are mirrored in RAID 1 so that a single 300 GB disk appears available to the operating system installer.

Procedure

1. Install Red Hat Enterprise Linux Version 7.1 or later, according to the manufacturer instructions. Obtain a bootable DVD that contains Red Hat Enterprise Linux Version 7.1 and start your system from this DVD. See the following guidance for installation options. If an item is not mentioned in the following list, leave the default selection.
 - a. After you start the DVD, choose Install or upgrade an existing system from the menu.
 - b. On the Welcome screen, select Test this media & install Red Hat Enterprise Linux 7.1.
 - c. Select your language and keyboard preferences.
 - d. Select your location to set the correct time zone.
 - e. Select Software Selection and then on the next screen, select Server with GUI.
 - f. From the installation summary page, click Installation Destination and verify the following items:
 - The local 300 GB disk is selected as the installation target.
 - Under Other Storage Options, Automatically configure partitioning is selected.
Click Done.
 - g. Click Begin Installation. After the installation starts, set the root password for your root user account.

After the installation is completed, restart the system and log in as the root user. Issue the `df` command to verify your basic partitioning. For example, on a test system, the initial partitioning produced the following result:

```
[root@tvapp02]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  50G   3.0G   48G   6% /
devtmpfs        32G     0   32G   0% /dev
tmpfs           32G   92K   32G   1% /dev/shm
tmpfs           32G   8.8M   32G   1% /run
tmpfs           32G     0   32G   0% /sys/fs/cgroup
/dev/mapper/rhel-home 220G   37M  220G   1% /home
/dev/sda1       497M  124M  373M  25% /boot
```

2. Configure your TCP/IP settings according to the operating system installation instructions.

For optimal throughput and reliability, consider bonding multiple network ports together. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection, which aggregates several subordinate ports into a single logical connection. The preferred method is to use a bond mode of 802.3ad, miimon setting of 100, and a `xmit_hash_policy` setting of layer3+4.

Restriction: To use an LACP network connection, you must have a network switch that supports LACP.

For additional instructions about configuring bonded network connections with Red Hat Enterprise Linux Version 7, see [Create a Channel Bonding Interface](#).

3. Open the `/etc/hosts` file and complete the following actions:
 - o Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7  server.yourdomain.com  server
```

- o Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1 localhost
```

4. Install components that are required for the server installation. Complete the following steps to create a Yellowdog Updater Modified (YUM) repository and install the prerequisite packages.

- a. Mount your Red Hat Enterprise Linux installation DVD to a system directory. For example, to mount it to the /mnt directory, issue the following command:

```
mount -t iso9660 -o ro /dev/cdrom /mnt
```

- b. Verify that the DVD mounted by issuing the mount command. You should see output similar to the following example:

```
/dev/sr0 on /mnt type iso9660
```

- c. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

If the repos.d directory does not exist, create it.

- d. List directory contents:

```
ls rhel-source.repo
```

- e. Rename the original repo file by issuing the mv command. For example:

```
mv rhel-source.repo rhel-source.repo.orig
```

- f. Create a new repo file by using a text editor. For example, to use the vi editor, issue the following command:

```
vi rhel71_dvd.repo
```

- g. Add the following lines to the new repo file. The baseurl parameter specifies your directory mount point:

```
[rhel71_dvd]
name=DVD Redhat Enterprise Linux 7.1
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

- h. Install the prerequisite package ksh.x86_64, by issuing the yum command. For example:

```
yum install ksh.x86_64
```

Exception: You do not need to install the compat-libstdc++-33-3.2.3-69.el6.i686 and libstdc++.i686 libraries for Red Hat Enterprise Linux Version 7.1.

5. When the software installation is complete, you can restore the original YUM repository values by completing the following steps:

- a. Unmount the Red Hat Enterprise Linux installation DVD by issuing the following command:

```
umount /mnt
```

- b. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

- c. Rename the repo file that you created:

```
mv rhel71_dvd.repo rhel71_dvd.repo.orig
```

- d. Rename the original file to the original name:

```
mv rhel-source.repo.orig rhel-source.repo
```

6. Determine whether kernel parameter changes are required. Complete the following steps:

- a. Use the sysctl -a command to list the parameter values.
- b. Analyze the results by using the guidelines in Table 1 to determine whether any changes are required.
- c. If changes are required, set the parameters in the /etc/sysctl.conf file. The file changes are applied when the system is started.

Tip: Automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings. On Linux, the DB2® database software automatically adjusts interprocess communication (IPC) kernel parameter values to the preferred settings. For more information about kernel parameter settings, search for Linux kernel parameters in the IBM DB2 Version 11.1 product documentation.

Table 1. Linux kernel parameter optimum settings

Parameter	Description
kernel.shmni	The maximum number of segments.
kernel.shmmax	The maximum size of a shared memory segment (bytes). This parameter must be set before automatically starting the IBM Spectrum Protect™ server on system startup.
kernel.shmall	The maximum allocation of shared memory pages (pages).
kernel.sem	(SEMMSL) The maximum semaphores per array.
There are four values for the kernel.sem parameter.	(SEMMNS) The maximum semaphores per system.
	(SEMOPM) The maximum operations per semaphore call.
	(SEMMNI) The maximum number of arrays.
kernel.msgmni	The maximum number of system-wide message queues.
kernel.msgmax	The maximum size of messages (bytes).
kernel.msgmnb	The default maximum size of queue (bytes).
kernel.randomize_va_space	The kernel.randomize_va_space parameter configures the use of memory ASLR for the kernel. Disable ASLR because it can cause errors for the DB2 software. To learn more details about the Linux ASLR and DB2, see technote 1365583.
vm.swappiness	The vm.swappiness parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the DB2 product information.
vm.overcommit_memory	The vm.overcommit_memory parameter influences how much virtual memory the kernel permits allocating. For more information about kernel parameters, see the DB2 product information.

7. Open firewall ports to communicate with the server. Complete the following steps:

- a. Determine the zone that is used by the network interface. The zone is public, by default.

Issue the following command:

```
# firewall-cmd --get-active-zones
public
interfaces: ens4f0
```

- b. To use the default port address for communications with the server, open TCP/IP port 1500 in the Linux firewall.

Issue the following command:

```
firewall-cmd --zone=public --add-port=1500/tcp --permanent
```

If you want to use a value other than the default, you can specify a number in the range 1024 - 32767. If you open a port other than the default, you will need to specify that port when you run the configuration script.

- c. If you plan to use this system as a hub, open port 11090, which is the default port for secure (https) communications.

Issue the following command:

```
firewall-cmd --zone=public --add-port=11090/tcp --permanent
```

d. Reload the firewall definitions for the changes to take effect.

Issue the following command:

```
firewall-cmd --reload
```

8. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in Table 2. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 2. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	FSIZE	Unlimited	ulimit -Hf
Maximum number of open files	nfile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Windows systems

Install Microsoft Windows Server 2012 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Spectrum Protect™ server.

Procedure

1. Install Windows Server 2016 Standard Edition, according to the manufacturer instructions.
2. Change the Windows account control policies by completing the following steps.
 - a. Open the Local Security Policy editor by running secpol.msc.
 - b. Click Local Policies > Security Options and ensure that the following User Account Control policies are disabled:
 - Admin Approval Mode for the Built-in Administrator account
 - Run all administrators in Admin Approval Mode
3. Configure your TCP/IP settings according to installation instructions for the operating system.
4. Apply Windows updates and enable optional features by completing the following steps:
 - a. Apply the latest Windows Server 2016 updates.
 - b. Install and enable the Windows 2012 R2 feature Microsoft .NET Framework 3.5 from the Windows Server Manager.
 - c. If required, update the FC and Ethernet HBA device drivers to newer levels.
 - d. Install the multipath I/O driver that is appropriate for the disk system that you are using.
5. Open the default TCP/IP port, 1500, for communications with the IBM Spectrum Protect server. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Backup server port 1500"  
dir=in action=allow protocol=TCP localport=1500
```

6. On the Operations Center hub server, open the default port for secure (https) communications with the Operations Center. The port number is 11090. For example, issue the following command:

```
netsh advfirewall firewall add rule name="Operations Center port 11090"  
dir=in action=allow protocol=TCP localport=11090
```

Configuring multipath I/O

You can enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.

- AIX systems
Complete the following steps to enable and configure multipathing for disk storage.
- Linux systems
Complete the following steps to enable and configure multipathing for disk storage.
- Windows systems
Complete the following steps to enable and configure multipathing for disk storage.

AIX systems

Complete the following steps to enable and configure multipathing for disk storage.

Procedure

1. Determine the Fibre Channel port address that you must use for the host definition on the disk subsystem. Issue the `lscfg` command for every port.

- On small and medium systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
```

- On large systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
lscfg -vps -l fcs2 | grep "Network Address"
lscfg -vps -l fcs3 | grep "Network Address"
```

2. Ensure that the following AIX® file sets are installed:

- `devices.common.IBM.mpio.rte`
- `devices.fcp.disk.array.rte`
- `devices.fcp.disk.rte`

3. Issue the `cfgmgr` command to have AIX rescan the hardware and discover available disks. For example:

```
cfgmgr
```

4. To list the available disks, issue the following command:

```
lsdev -Ccdisk
```

You should see output similar to the following:

```
hdisk0 Available 00-00-00 SAS Disk Drive
hdisk1 Available 00-00-00 SAS Disk Drive
hdisk2 Available 01-00-00 SAS Disk Drive
hdisk3 Available 01-00-00 SAS Disk Drive
hdisk4 Available 06-01-02 MPIO IBM 2076 FC Disk
hdisk5 Available 07-01-02 MPIO IBM 2076 FC Disk
...
```

5. Use the output from the `lsdev` command to identify and list device IDs for each disk device.

For example, a device ID could be `hdisk4`. Save the list of device IDs to use when you create file systems for the IBM Spectrum Protect™ server.

6. Correlate the SCSI device IDs to specific disk LUNs from the disk system by listing detailed information about all physical volumes in the system. Issue the following command:

```
lspv -u
```

On an IBM® Storwize® system, the following information is an example of what is shown for each device:

```
hdisk4 00f8cf083fd97327 None active
332136005076300810105780000000000000003004214503IBMfcp
```

In the example, `60050763008101057800000000000030` is the UID for the volume, as reported by the Storwize management interface.

To verify disk size in megabytes and compare the value with what is listed for the system, issue the following command:

```
bootinfo -s hdisk4
```

Linux systems

Complete the following steps to enable and configure multipathing for disk storage.

Procedure

1. Edit the `/etc/multipath.conf` file to enable multipathing for Linux hosts. If the `multipath.conf` file does not exist, you can create it by issuing the following command:

```
mpathconf --enable
```

The following parameters were set in `multipath.conf` for testing on an IBM Storwize® system:

```
defaults {
    user_friendly_names no
}

devices {
    device {
        vendor "IBM "
        product "2145"
        path_grouping_policy group_by_prio
        user_friendly_names no
        path_selector "round-robin 0"
        prio "alua"
        path_checker "tur"
        failback "immediate"
        no_path_retry 5
        rr_weight uniform
        rr_min_io_rq "1"
        dev_loss_tmo 120
    }
}
```

2. Set the multipath option to start when the system is started. Issue the following commands:

```
systemctl enable multipathd.service
systemctl start multipathd.service
```

3. To verify that disks are visible to the operating system and are managed by multipath, issue the following command:

```
multipath -l
```

4. Ensure that each device is listed and that it has as many paths as you expect. You can use size and device ID information to identify which disks are listed.

For example, the following output shows that a 2 TB disk has two path groups and four active paths. The 2 TB size confirms that the disk corresponds to a pool file system. Use part of the long device ID number (12, in this example) to search for the volume on the disk-system management interface.

```
[root@tapsrv01 code]# multipath -l
36005076802810c509800000000000012 dm-43 IBM,2145
size=2.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=0 status=active
| |- 2:0:1:18 sdcw 70:64 active undef running
| `-- 4:0:0:18 sdgb 131:112 active undef running
`+- policy='round-robin 0' prio=0 status=enabled
  |- 1:0:1:18 sdat 66:208 active undef running
  `-- 3:0:0:18 sddy 128:0 active undef running
```

- a. If needed, correct disk LUN host assignments and force a bus rescan. For example:

```
echo "- - -" > /sys/class/scsi_host/host0/scan
echo "- - -" > /sys/class/scsi_host/host1/scan
echo "- - -" > /sys/class/scsi_host/host2/scan
```

You can also restart the system to rescan disk LUN host assignments.

- b. Confirm that disks are now available for multipath I/O by reissuing the `multipath -l` command.

5. Use the multipath output to identify and list device IDs for each disk device.

For example, the device ID for your 2 TB disk is 36005076802810c50980000000000012.

Save the list of device IDs to use in the next step.

Windows systems

Complete the following steps to enable and configure multipathing for disk storage.

Procedure

1. Ensure that the Multipath I/O feature is installed. If needed, install additional vendor-specific multipath drivers.
2. To verify that disks are visible to the operating system and are managed by multipath I/O, issue the following command:

```
c:\program files\IBM\SDDDSM\datapath.exe query device
```

3. Review the multipath output and ensure that each device is listed and that it has as many paths as you expect. You can use size and device serial information to identify which disks are listed.

For example, by using part of the long device serial number (34, in this example) you can search for the volume on the disk-system management interface. The 2 TB size confirms that the disk corresponds to a storage pool file system.

```
DEV#: 4 DEVICE NAME: Disk5 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 60050763008101057800000000000034 LUN SIZE: 2.0TB
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 0 0
1 Scsi Port2 Bus0/Disk5 Part0 OPEN NORMAL 27176 0
2 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 28494 0
3 Scsi Port3 Bus0/Disk5 Part0 OPEN NORMAL 0 0
```

4. Create a list of disk device IDs by using the serial numbers that are returned from the multipath output in the previous step.

For example, the device ID for your 2 TB disk is 60050763008101057800000000000034

Save the list of device IDs to use in the next step.

5. To bring new disks online and clear the read-only attribute, run diskpart.exe with the following commands. Repeat for each of the disks:

```
diskpart
select Disk 1
online disk
attribute disk clear readonly
select Disk 2
online disk
attribute disk clear readonly
< ... >
select Disk 49
online disk
attribute disk clear readonly
exit
```

Creating the user ID for the server

Create the user ID that owns the IBM Spectrum Protect™ server instance. You specify this user ID when you create the server instance during initial configuration of the server.

About this task

You can specify only lowercase letters (a-z), numerals (0-9), and the underscore character (_) for the user ID. The user ID and group name must comply with the following rules:

- The length must be 8 characters or fewer.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

Procedure

1. Use operating system commands to create a user ID.

- o **AIX** | **Linux** Create a group and user ID in the home directory of the user that owns the server instance.

For example, to create the user ID `tsminst1` in group `tsmsrvrs` with a password of `tsminst1`, issue the following commands from an administrative user ID:

AIX

```
mkgroup id=1001 tsmsrvrs
mkuser id=1002 pgrp=tsmsrvrs home=/home/tsminst1 tsminst1
passwd tsminst1
```

Linux

```
groupadd tsmsrvrs
useradd -d /home/tsminst1 -m -g tsmsrvrs -s /bin/bash tsminst1
passwd tsminst1
```

Log off, and then log in to your system. Change to the user account that you created. Use an interactive login program, such as `telnet`, so that you are prompted for the password and can change it if necessary.

- o **Windows** Create a user ID and then add the new ID to the Administrators group. For example, to create the user ID `tsminst1`, issue the following command:

```
net user tsminst1 * /add
```

After you create and verify a password for the new user, add the user ID to the Administrators group by issuing the following commands:

```
net localgroup Administrators tsminst1 /add
net localgroup DB2ADMNS tsminst1 /add
```

2. Log off the new user ID.

Preparing file systems for the server

You must complete file system configuration for the disk storage to be used by the server.

- Preparing file systems on AIX systems
You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.
- Preparing file systems on Linux systems
You must format `ext4` or `xfs` file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.
- Preparing file systems on Windows systems
You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect server.

Preparing file systems on AIX systems

You must create volume groups, logical volumes, and file systems for the server by using the AIX® Logical Volume Manager.

Procedure

1. Increase the queue depth and maximum transfer size for all of the available `hdiskX` disks. Issue the following commands for each disk:

```
chdev -l hdisk4 -a max_transfer=0x100000
chdev -l hdisk4 -a queue_depth=32
chdev -l hdisk4 -a reserve_policy=no_reserve
chdev -l hdisk4 -a algorithm=round_robin
```

Do not run these commands for operating system internal disks, for example, `hdisk0`.

2. Create volume groups for the IBM Spectrum Protect™ database, active log, archive log, database backup, and storage pool. Issue the `mkvg` command, specifying the device IDs for corresponding disks that you previously identified. For example, if the device names `hdisk4`, `hdisk5`, and `hdisk6` correspond to database disks, include them in the database volume group and so on.

System size: The following commands are based on the medium system configuration. For small and large systems, you must adjust the syntax as required.

```
mkvg -S -y tsmdb hdisk2 hdisk3 hdisk4
mkvg -S -y tsmactlog hdisk5
mkvg -S -y tsmarchlog hdisk6
mkvg -S -y tsmdbback hdisk7 hdisk8 hdisk9 hdisk10
mkvg -S -y tsmstgpool hdisk11 hdisk12 hdisk13 hdisk14 ... hdisk49
```

3. Determine the physical volume names and the number of free physical partitions to use when you create logical volumes. Issue the `lsvg` for each volume group that you created in the previous step.

For example:

```
lsvg -p tsmdb
```

The output is similar to the following. The *FREE PPs* column represents the free physical partitions:

```
tsmdb:
PV_NAME  PV STATE  TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk4   active    1631      1631      327..326..326..326..326
hdisk5   active    1631      1631      327..326..326..326..326
hdisk6   active    1631      1631      327..326..326..326..326
```

4. Create logical volumes in each volume group by using the `mklv` command. The volume size, volume group, and device names vary, depending on the size of your system and variations in your disk configuration.

For example, to create the volumes for the IBM Spectrum Protect database on a medium system, issue the following commands:

```
mklv -y tsmdb00 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk2
mklv -y tsmdb01 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk3
mklv -y tsmdb02 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk4
```

5. Format file systems in each logical volume by using the `crfs` command.

For example, to format file systems for the database on a medium system, issue the following commands:

```
crfs -v jfs2 -d tsmdb00 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace00 -A yes
crfs -v jfs2 -d tsmdb01 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace01 -A yes
crfs -v jfs2 -d tsmdb02 -p rw -a logname=INLINE -a options=rbrw
-a agblksize=4096 -m /tsminst1/TSMdbspace02 -A yes
```

6. Mount all of the newly created file systems by issuing the following command:

```
mount -a
```

7. List all file systems by issuing the `df` command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example of command output shows that the amount of used space is typically 1%:

```
tapsrv07> df -g /tsminst1/*
Filesystem      GB blocks   Free    %Used   Iused   %Iused  Mounted on
/dev/tsmact00   195.12    194.59    1%      4        1%     /tsminst1/TSMalog
```

8. Verify that the user ID that you created in Creating the user ID for the server has read and write access to the directories for the server.

Preparing file systems on Linux systems

You must format ext4 or xfs file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Using the list of device IDs that you generated previously, issue the `mkfs` command to create and format a file system for each storage LUN device. Specify the device ID in the command. See the following examples. For the database, format ext4 file systems:

```
mkfs -t ext4 -T largefile -m 2 /dev/mapper/36005076802810c50980000000000012
```

For storage pool LUNs, format xfs file systems:

```
mkfs -t xfs /dev/mapper/3600507630081010578000000000002c3
```

You might issue the mkfs command as many as 50 times, depending on how many different devices you have.

2. Create mount point directories for file systems.

Issue the mkdir command for each directory that you must create. Use the directory values that you recorded in the planning worksheets.

For example, to create the server instance directory by using the default value, issue the following command:

```
mkdir /tsminst1
```

Repeat the mkdir command for each file system.

3. Add an entry in the /etc/fstab file for each file system so that file systems are mounted automatically when the server is started.

For example:

```
/dev/mapper/36005076802810c509800000000000012 /tsminst1/TSMdbspace00 ext4 defaults 0 0
```

4. Mount the file systems that you added to the /etc/fstab file by issuing the mount -a command.

5. List all file systems by issuing the df command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The following example on an IBM® Storwize® system shows that the amount of used space is typically 1%:

```
[root@tapsrv04 ~]# df -h /tsminst1/*
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/36005076300810105780000000000003 134G  188M 132G  1%  /tsminst1/TSMalog
```

6. Verify that the user ID that you created in Creating the user ID for the server has read and write access to the directories for the IBM Spectrum Protect server.

Preparing file systems on Windows systems

You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Spectrum Protect™ server.

Procedure

1. Create mount point directories for file systems.

Issue the md command for each directory that you must create. Use the directory values that you recorded in the planning worksheets. For example, to create the server instance directory by using the default value, issue the following command:

```
md c:\tsminst1
```

Repeat the md command for each file system.

2. Create a volume for every disk LUN that is mapped to a directory under the server instance directory by using the Windows volume manager.

Go to Server Manager > File and Storage Services and complete the following steps for each disk that corresponds to the LUN mapping that was created in the previous step:

a. Bring the disk online.

b. Initialize the disk to the GPT basic type, which is the default.

c. Create a simple volume that occupies all of the space on the disk. Format the file system by using NTFS, and assign a label that matches the purpose of the volume, such as TSMfile00. Do not assign the new volume to a drive letter. Instead, map the volume to a directory under the instance directory, such as C:\tsminst1\TSMfile00.

Tip: Determine the volume label and directory mapping labels based on the size of the disk that is reported.

3. Verify that file systems are mounted at the correct LUN and correct mount point. List all file systems by issuing the mountvol command and then review the output. For example:

```
\\?\Volume{8ffb9678-3216-474c-a021-20e420816a92}\
C:\tsminst1\TSMdbspace00\
```

4. After the disk configuration is complete, restart the system.

What to do next

You can confirm the amount of free space for each volume by using Windows Explorer.

Installing the server and Operations Center

Use the IBM® Installation Manager graphical wizard to install the components.

- Installing on AIX and Linux systems
Install the IBM Spectrum Protect™ server and the Operations Center on the same system.
- Installing on Windows systems
Install the IBM Spectrum Protect server and the Operations Center on the same system.

Installing on AIX® and Linux systems

Install the IBM Spectrum Protect™ server and the Operations Center on the same system.

Before you begin

Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

1. **AIX** Verify that the required RPM files are installed on your system.

See Installing prerequisite RPM files for the graphical wizard for details.

2. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042992.
3. Go to Passport Advantage® and download the package file to an empty directory of your choice.
4. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

5. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file.

6. **AIX** Ensure that the following command is enabled so that the wizards work properly:

```
lsuser
```

By default, the command is enabled.

7. Change to the directory where you placed the executable file.
8. Start the installation wizard by issuing the following command:

```
./install.sh
```

When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.

- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.

- Installing prerequisite RPM files for the graphical wizard
RPM files are required for the IBM Installation Manager graphical wizard.

Installing on Windows systems

Install the IBM Spectrum Protect™ server and the Operations Center on the same system.

Before you begin

Make sure that the following prerequisites are met:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

1. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document at technote 4042993.
2. Go to Passport Advantage® and download the package file to an empty directory of your choice.
3. Change to the directory where you placed the executable file.
4. Double-click the executable file to extract to the current directory.
5. In the directory where the installation files were extracted, start the installation wizard by double-clicking the install.bat file. When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory.

To view installation log files from the Installation Manager tool, click File > View Log. To collect these log files from the Installation Manager tool, click Help > Export Data for Problem Analysis.

- After you install the server and before you customize it for your use, go to the IBM Spectrum Protect support site. Click Support and downloads and apply any applicable fixes.

Configuring the server and the Operations Center

After you install the components, complete the configuration for the IBM Spectrum Protect™ server and the Operations Center.

- **Configuring the server instance**
Use the IBM Spectrum Protect server instance configuration wizard to complete the initial configuration of the server.
- **Installing the backup-archive client**
As a best practice, install the IBM Spectrum Protect backup-archive client on the server system so that the administrative command-line client and scheduler are available.
- **Setting options for the server**
Review the server options file that is installed with the IBM Spectrum Protect server to verify that the correct values are set for your system.
- **Security concepts**
You can protect IBM Spectrum Protect from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.
- **Configuring the Operations Center**
After you install the Operations Center, complete the following configuration steps to start managing your storage environment.
- **Registering the product license**
To register your license for the IBM Spectrum Protect product, use the REGISTER LICENSE command.
- **Defining data retention rules for your business**
After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.
- **Defining schedules for server maintenance activities**
Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations

- Center command builder.
- Defining client schedules
Use the Operations Center to create schedules for client operations.

Configuring the server instance

Use the IBM Spectrum Protect™ server instance configuration wizard to complete the initial configuration of the server.

Before you begin

Ensure that the following requirements are met:

AIX | **Linux**

- The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/directory`. Also, ensure that the SSH daemon service has access rights to connect to the system by using the `localhost` value.
- You must be able to log in to IBM Spectrum Protect with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

Windows

Verify that the remote registry service is started by completing the following steps:

- Click Start > Administrative Tools > Services. In the Services window, select Remote Registry. If it is not started, click Start.
- Ensure that port 137, 139, and 445 are not blocked by a firewall:
 - Click Start > Control Panel > Windows Firewall.
 - Select Advanced Settings.
 - Select Inbound Rules.
 - Select New Rule.
 - Create a port rule for TCP ports 137, 139, and 445 to allow connections for domain and private networks.
- Configure the user account control by accessing the local security policy options and completing the following steps.
 - Click Start > Administrative Tools > Local Security Policy. Expand Local Policies > Security Options.
 - If not already enabled, enable the built-in administrator account by selecting Accounts: Administrator account status > Enable > OK.
 - If not already disabled, disable user account control for all Windows administrators by selecting User Account Control: Run all administrators in Admin Approval Mode > Disable > OK.
 - If not already disabled, disable the User Account Control for the built-in Administrator account by selecting User Account Control: Admin Approval Mode for the Built-in Administrator Account > Disable > OK.
- If you changed any settings in the preceding steps, restart the server before you proceed with the configuration wizard.

About this task

The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Procedure

- Start the local version of the wizard.
 - AIX** | **Linux** Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be only run as a root user.
 - Windows** Click Start > All Programs > IBM Spectrum Protect > Configuration Wizard.
- Follow the instructions to complete the configuration. Use the information that you recorded in Planning worksheets during IBM Spectrum Protect system setup to specify directories and options in the wizard.

AIX | **Linux**

On the Server Information window, set the server to start automatically by using the instance user ID when the system boots.

Windows

By using the configuration wizard, the server is set to start automatically when rebooted.

Installing the backup-archive client

As a best practice, install the IBM Spectrum Protect™ backup-archive client on the server system so that the administrative command-line client and scheduler are available.

Procedure

To install the backup-archive client, follow the installation instructions for your operating system.

- Install UNIX and Linux backup-archive clients
- Installing the Windows client for the first time

Setting options for the server

Review the server options file that is installed with the IBM Spectrum Protect™ server to verify that the correct values are set for your system.

Procedure

1. Go to the server instance directory and open the dsmserv.opt file.
2. Review the values in the following table and verify your server option settings, based on system size.

Server option	Value
ACTIVELOGDIRECTORY	Directory path that was specified during configuration
ACTIVELOGSIZE	131072
ARCHLOGCOMPRESS	No
ARCHLOGDIRECTORY	Directory path that was specified during configuration
COMMMETHOD	TCPIP
COMMTIMEOUT	3600
DEVCONFIG	devconf.dat
EXPINTERVAL	0
IDLETIMEOUT	60
MAXSESSIONS	500
NUMOPENVOLSAALLOWED	20
TCPADMINPORT	1500
TCPPORT	1500
VOLUMEHISTORY	volhist.dat

Update server option settings if necessary, to match the values in the table. To make updates, close the dsmserv.opt file and use the SETOPT command from the administrative command-line interface to set the options.

For example, to update the IDLETIMEOUT option to 60, issue the following command:

```
setopt idletimeout 60
```

3. To configure secure communications for the server, clients, and the Operations Center, verify the options in the following table.

Server option	All system sizes
SSLDISABLELEGACYTLS	YES
SSLFIPSMODE	NO
SSLTCPPORT	Specify the SSL port number. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.
SSLTCPADMINPORT	Specify the port address on which the server waits for requests for SSL-enabled sessions from the command-line administrative client.

Server option	All system sizes
SSLTLS12	YES

If any of the option values must be updated, edit the dsmserv.opt file by using the following guidelines:

- Remove the asterisk at the beginning of a line to enable an option.
- On each line, enter only one option and the specified value for the option.
- If an option occurs in multiple entries in the file, the server uses the last entry.

Save your changes and close the file. If you edit the dsmserv.opt file directly, you must restart the server for the changes to take effect.

Security concepts

You can protect IBM Spectrum Protect™ from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.

Transport Layer Security

You can use the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol to provide transport layer security for a secure connection between servers, clients, and storage agents. If you send data between the server, client, and storage agent, use SSL or TLS to encrypt the data.

Tip: Any IBM Spectrum Protect documentation that indicates "SSL" or to "select SSL" applies to TLS.

SSL is provided by the Global Security Kit (GSKit) that is installed with the IBM Spectrum Protect server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with a DB2® database instance that is used by any IBM Spectrum Protect servers.

Each server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a certificate authority (CA). You can use your own certificates or purchase certificates from a CA. Either certificate must be installed and added to the key database on the IBM Spectrum Protect server, client, or storage agent. The certificate is verified by the SSL client or server that requests or initiates the SSL communication. Some CA certificates are preinstalled in the key databases, by default.

SSL is set up independently on the IBM Spectrum Protect server, client, and storage agent.

Authority levels

With each IBM Spectrum Protect server, different administrative authority levels are available that determine the tasks that an administrator can complete.

After registration, an administrator must be granted authority by being assigned one or more administrative authority levels. An administrator with system authority can complete any task with the server and assign authority levels to other administrators by using the GRANT AUTHORITY command. Administrators with policy, storage, or operator authority can complete subsets of tasks.

An administrator can register other administrator IDs, grant levels of authority to them, rename IDs, remove IDs, and lock and unlock them from the server.

An administrator can control access to specific client nodes for root user IDs and non-root user IDs. By default, a non-root user ID cannot back up data on the node. Use the UPDATE NODE command to change the node settings to enable backup.

Passwords

By default, the server automatically uses password authentication. With password authentication, all users must enter a password when they access the server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see Managing passwords and logon procedures (V7.1.1).

Table 1. Password authentication characteristics

Characteristic	More information
----------------	------------------

Characteristic	More information
Case-sensitivity	Not case-sensitive.
Default password expiration	90 days. The expiration period begins when an administrator ID or client node is first registered to the server. If the password is not changed within this period, the password must be changed the next time that the user accesses the server.
Invalid password attempts	You can set a limit on consecutive invalid password attempts for all client nodes. When the limit is exceeded, the server locks the node.
Default password length	8 characters. The administrator can specify a minimum length. Beginning with Version 8.1.4, the default minimum length for server passwords changed from 0 to 8 characters.

Session security

Session security is the level of security that is used for communication among IBM Spectrum Protect client nodes, administrative clients, and servers and is set by using the SESSIONSECURITY parameter.

The SESSIONSECURITY parameter can be set to one of the following values:

- The STRICT value enforces the highest level of security for communication between IBM Spectrum Protect servers, nodes, and administrators.
- The TRANSITIONAL value specifies that the existing communication protocol is used while you update your IBM Spectrum Protect software to V8.1.2 or later. This is the default. When SESSIONSECURITY=TRANSITIONAL, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

Note: You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Spectrum Protect servers, but you are not required to upgrade servers first. Communication between servers and clients is not interrupted.

For more information about the SESSIONSECURITY parameter values, see the following commands.

Table 2. Commands used to set the SESSIONSECURITY parameter

Entity	Command
Client nodes	<ul style="list-style-type: none"> • REGISTER NODE • UPDATE NODE
Administrators	<ul style="list-style-type: none"> • REGISTER ADMIN • UPDATE ADMIN
Servers	<ul style="list-style-type: none"> • DEFINE SERVER • UPDATE SERVER

Administrators that authenticate by using the DSMADMC command, DSMC command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

Tips:

- Ensure that all IBM Spectrum Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.

- After an administrator authenticates to a V8.1.2 or later server by using a V8.1.2 or later client, the administrator can authenticate only on clients or servers that are using V8.1.2 or later. An administrator command can be issued from any system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Enforce the highest level of security for communication with the IBM Spectrum Protect server by ensuring that all nodes, administrators, and servers use STRICT session security. You can use the SELECT command to determine which servers, nodes, and administrators are using TRANSITIONAL session security and should be updated to use STRICT session security.

- **Configuring secure communications with Transport Layer Security**
To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) is enabled on the IBM Spectrum Protect server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.

Related tasks:

[🔗 Securing communications](#)

Configuring the Operations Center

After you install the Operations Center, complete the following configuration steps to start managing your storage environment.

Before you begin

When you connect to the Operations Center for the first time, you must provide the following information:

- Connection information for the server that you want to designate as a hub server
- Login credentials for an administrator ID that is defined for that server

Procedure

1. Designate the hub server. In a web browser, enter the following address:

```
https://hostname:secure_port/oc
```

where:

- *hostname* represents the name of the computer where the Operations Center is installed
- *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer

For example, if your host name is tsm.storage.mylocation.com and you are using the default secure port for the Operations Center, which is 11090, the address is:

```
https://tsm.storage.mylocation.com:11090/oc
```

When you log in to the Operations Center for the first time, a wizard guides you through an initial configuration to set up a new administrator with system authority on the server.

2. Set up secure communications between the Operations Center and the hub server by configuring the Secure Sockets Layer (SSL) protocol.

Follow the instructions in [Securing communications between the Operations Center and the hub server](#).

3. Optional: To receive a daily email report that summarizes system status, configure your email settings in the Operations Center.

Follow the instructions in [Tracking system status by using email reports](#).

- **Securing communications between the Operations Center and the hub server**
To secure communications between the Operations Center and the hub server, add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.

Registering the product license


To register your license for the IBM Spectrum Protect™ product, use the REGISTER LICENSE command.

About this task

Licenses are stored in enrollment certificate files, which contain licensing information for the product. The enrollment certificate files are on the installation media, and are placed on the server during installation. When you register the product, the licenses are stored in a NODELOCK file within the current directory.

Procedure


Register a license by specifying the name of the enrollment certificate file that contains the license. To use the Operations Center command builder for this task, complete the following steps.

1. Open the Operations Center.
2. Open the Operations Center command builder by hovering over the settings icon  and clicking Command Builder.
3. Issue the REGISTER LICENSE command. For example, to register a base IBM Spectrum Protect license, issue the following command:

```
register license file=tsmbasic.lic
```

What to do next

Save the installation media that contains your enrollment certificate files. You might need to register your license again if, for example, one of the following conditions occur:

- The server is moved to a different computer.
- The NODELOCK file is corrupted. The server stores license information in the NODELOCK file, which is in the directory from which the server is started.
-  If you change the processor chip that is associated with the server on which the server is installed.

Defining data retention rules for your business

After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.

Procedure

1. On the Services page of the Operations Center, select the STANDARD domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab. The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.
3. Click the Configure toggle, and make the following changes:
 - Change the backup destination for the STANDARD management class to the directory-container storage pool.
 - Change the value for the Backups column to No limit.
 - Change the retention period. Set the Keep Extra Backups column to 30 days or more, depending on your business requirements.
4. Save your changes and click the Configure toggle again so that the policy set is no longer editable.
5. Activate the policy set by clicking Activate.


Defining schedules for server maintenance activities

Create schedules for each server maintenance operation by using the DEFINE SCHEDULE command in the Operations Center command builder.

About this task

Schedule server maintenance operations to run after client backup operations. You can control the timing of schedules by setting the start time in combination with the duration time for each operation.

The following figure provides an example of how to plan maintenance operations.
Figure 1. Daily schedule of server operations for a tape solution



The following table shows how you can schedule server maintenance processes in combination with the client backup schedule for a tape solution.

Operation	Schedule
Client backup	Starts at 11 PM.
Storage pool backup	Starts at 6 AM.
Processing for database and disaster recovery files	<ul style="list-style-type: none"> The database backup operation starts at 10 AM, or 11 hours after the beginning of the client backup operation. This process runs until completion. Device configuration information and volume history backup operations start at 5 PM, or 7 hours after the start of the database backup operation. Volume history deletion starts at 8 PM, or 10 hours after the start of the database backup operation.
Preparation of files for offsite vaulting	Starts at 10 AM, at the same time as processing for the database and disaster recovery files.
Migration from disk to tape	Starts at 12 PM, or 2 hours after the start of the database backup operation.
Inventory expiration	Starts at 2 PM, or 15 hours after the beginning of the client backup operation. This process runs until completion.
Space reclamation	Starts at 3 PM, or 16 hours after the beginning of the client backup operation.

Procedure

After you configure the device class for the database backup operations, create schedules for database backup and other required maintenance operations by using the DEFINE SCHEDULE command. Depending on the size of your environment, you might need to adjust the start times for each schedule in the example.

1. Define a device class for the backup operation before you create the schedule for database backups. Use the DEFINE DEVCLASS command to create a device class that is named LTOTAPE:

```
define devclass ltotape devtype=lto library=ltolib
```

2. Set the device class for automatic database backups. Use the SET DBRECOVERY command to specify the device class that you created for the database backup in the preceding step. For example, if the device class is LTOTAPE, issue the following command:

```
set dbrecovery ltotape
```

3. Create schedules for the maintenance operations by using the DEFINE SCHEDULE command. See the following table for the required operations with examples of the commands.

Operation	Example commands and additional information
Back up storage pools.	<p>Create a schedule to run the BACKUP STGPOOL command. For example, issue the following command to create a backup schedule for a primary storage pool that is named PRIMARY_POOL. The pool will be backed up to a copy storage pool, COPYSTG:</p> <pre>define schedule BACKUPSTGPOOL type=administrative cmd="backup stgpool primary_pool copystg" active=yes starttime=06:00 period=1</pre>
Back up the database.	<p>Create a schedule to run the BACKUP DB command. For example, issue the following command to create a backup schedule that uses the new device class:</p> <pre>define schedule DBBACKUP type=admin cmd="backup db devclass=ltotape type=full numstreams=3 wait=yes compress=yes" active=yes desc="Back up the database." startdate=today starttime=10:00:00 duration=45 durunits=minutes</pre>
Replicate nodes.	<p>Optionally, use node replication to protect client data by backing the data up to a secondary server. For instructions, see Replicating client data to another server. Ensure that node replication is completed before migration operations begin.</p>

Operation	Example commands and additional information
Migrate data from disk to tape daily.	<p>Create a schedule for storage pool migration.</p> <p>For example, if a disk storage pool is named DISKPOOL and the next storage pool is TAPEPOOL, you can schedule storage pool migration by issuing the following command:</p> <pre>define schedule stgpool_migration type=administrative cmd="migrate stgpool diskpool lomig=0" active=yes description="migrate disk storagepool to tapepool" startdate=today starttime=12:00 duration=2 durunits=hours period=1 perunits=days</pre> <p>To maximize throughput, you can specify the number of parallel processes to use for migrating files by completing the following steps:</p> <ol style="list-style-type: none"> For the tape storage pool, ensure that collocation is enabled. To verify whether collocation is enabled, run the QUERY STGPOOL command. Verify that a value of GROUP, NODE, or FILESPACE is specified in the COLLOCATE field. If a value of GROUP, NODE, or FILESPACE is not specified, use the UPDATE STGPOOL command to specify COLLOCATE=GROUP, COLLOCATE=NODE, or COLLOCATE=FILESPACE, depending on your system configuration. For the disk storage pool, use the DEFINE STGPOOL or UPDATE STGPOOL command to specify a value for the MIGPROCESS parameter. For example, if you have 12 tape drives, specify MIGPROCESS=10. In this way, a maximum of 10 tape drives are used for migration processes. Two drives are reserved for other tasks, such as restore, database backup, and client backup operations.
Prepare files for offsite vaulting.	<ol style="list-style-type: none"> Move tape volumes offsite by following the instructions in Moving backup media. Create the disaster recovery plan file by issuing the PREPARE command on the source server: <pre>prepare</pre> Ensure that all volumes that are required for disaster recovery are included in the recovery plan file. For more information, see Preparing for and recovering from a disaster by using DRM.
Back up the device configuration information.	<p>Create a schedule to run the BACKUP DEVCONFIG command:</p> <pre>define schedule DEVCONFIGBKUP type=admin cmd="backup devconfig filenames=devconfig.dat" active=yes desc="Backup the device configuration file." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>
Back up the volume history.	<p>Create a schedule to run the BACKUP VOLHISTORY command:</p> <pre>define schedule VOLHISTBKUP type=admin cmd="backup volhistory filenames=volhist.dat" active=yes desc="Back up the volume history." startdate=today starttime=17:00:00 duration=45 durunits=minutes</pre>

Operation	Example commands and additional information
Remove older versions of database backups that are no longer required.	<p>Create a schedule to run the DELETE VOLHISTORY command:</p> <pre>define schedule DELVOLHIST type=admin cmd="delete volhistory type=dbb todate=today-6 totime=now" active=yes desc="Remove old database backups." startdate=today starttime=20:00:00 duration=45 durunits=minutes</pre>
Remove objects that exceed their allowed retention.	<p>Create a schedule to run the EXPIRE INVENTORY command.</p> <p>Set the RESOURCE parameter based on the system size that you are configuring to be equal to the number of processor cores that you specified for your system.</p> <p>For example, issue the following command to create a schedule that is named EXPINVENTORY:</p> <pre>define schedule EXPINVENTORY type=admin cmd="expire inventory wait=yes resource=8 duration=120" active=yes desc="Remove expired objects." startdate=today starttime=14:00:00 duration=1 durunits=hours</pre>
Reclaim space.	<p>Create a schedule to run the RECLAIM STGPOOL command.</p> <p>For example, issue the following command to create a schedule that is named RECLAIM:</p> <pre>define schedule RECLAIM type=admin cmd="reclaim stgpool tapepool duration=60" startdate=today starttime=15:00:00 duration=5 durunits=hours</pre> <p>Tip: To maximize throughput, you can specify the number of parallel processes to use for reclaiming space. Update the tape storage pool by using the UPDATE STGPOOL command and specify a value for the RECLAIMPROCESS parameter. For example, if you have 12 tape drives, specify RECLAIMPROCESS=5. Because two drives are used for each reclamation process, the total number of drives that can be used for reclamation is 10. Two drives are reserved for backup operations.</p>

What to do next

After you create schedules for the server maintenance tasks, you can view them in the Operations Center by completing the following steps:

1. On the Operations Center menu bar, hover over Servers.
2. Click Maintenance.

- Moving backup media
To recover from a disaster, you need database backup volumes, copy storage pool volumes, and additional files. To stay prepared for a disaster, you must complete daily tasks.

Related reference:

- [UPDATE STGPOOL \(Update a storage pool\)](#)
- [DEFINE SCHEDULE \(Define a schedule for an administrative command\)](#)

Related information:

- [DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Defining client schedules

Use the Operations Center to create schedules for client operations.

Procedure

1. On the Operations Center menu bar, hover over Clients.
2. Click Schedules.
3. Click +Schedule.
4. Complete the steps in the Create Schedule wizard. Set client backup schedules to start at 22:00, based on the server maintenance activities that you scheduled in Defining schedules for server maintenance activities.

Attaching tape devices for the server

Before the server can use a tape device, you must attach the device to your server system and install the appropriate tape device driver.

About this task

To optimize system performance, use fast, high-capacity tape devices. Provision enough tape drives to meet your business requirements.

Attach tape devices on their own host bus adapter (HBA), not shared with other devices types such as disk. IBM® tape drives have some special requirements for HBAs and associated drivers.

- **AIX** | **Linux** Attaching an automated library device to your system
You can attach an automated library device to your system to store your data on tapes.
- Selecting a tape device driver
To use tape devices with IBM Spectrum Protect™ you must install the appropriate tape device driver.
- **AIX** | **Linux** Special file names for tape devices
A special file name for a tape device is required for the IBM Spectrum Protect server to work with tape, medium changer, or removable media devices.
- Installing and configuring tape device drivers
Before you can use tape devices with IBM Spectrum Protect, you must install the correct tape device driver.

Attaching an automated library device to your system

You can attach an automated library device to your system to store your data on tapes.

About this task

Before you attach an automated library device, consider the following restrictions:

- Attached devices must be on their own Host Bus Adapter (HBA).
- An HBA must not be shared with other device types, such as a disk.
- For multiport Fibre Channel HBAs, devices must be attached on their own port. These ports must not be shared with other device types.
- IBM® tape drives have some special requirements on HBA and associated drivers. For more information about devices, see the website for your operating system:
 - IBM Spectrum Protect™ Supported Devices for AIX®
 - IBM Spectrum Protect Supported Devices for Linux and Windows

Procedure

To use the Fibre Channel (FC) adapter, complete the following steps:

1. Install the FC adapter and associated drivers.
 2. Install the appropriate device drivers for attached medium changer devices.
- **AIX** | **Linux** Setting the library mode
For the IBM Spectrum Protect server to access a SCSI library, the tape device must be set for the appropriate mode.

Related concepts:

Selecting a tape device driver

To use tape devices with IBM Spectrum Protect™ you must install the appropriate tape device driver.

- IBM tape device drivers
IBM® tape device drivers are available for most IBM labeled tape devices.
- IBM Spectrum Protect tape device drivers
The IBM Spectrum Protect server provides tape device drivers.

Related reference:

Installing and configuring tape device drivers

IBM tape device drivers

IBM® tape device drivers are available for most IBM labeled tape devices.

You can download IBM tape device drivers from the Fix Central website:

1. Go to the Fix Central website: Fix Central website.
2. Click Select product.
3. Select System Storage for the Product Group menu.
4. Select Tape systems for the System Storage menu.
5. Select Tape drivers and software for the Tape systems menu.
6. Select Tape device drivers for the Tape drivers and software menu. In addition to tape drivers, you also get access to tools such as the IBM Tape Diagnostic Tool (ITDT).
7. Select your operating system for the Platform menu.

AIX | Windows

For the most up-to-date list of devices and operating-system levels that are supported by IBM tape device drivers, see the IBM Spectrum Protect™ Supported Devices website at Supported devices for AIX and Windows.

Linux

For the most up-to-date list of tape devices and operating-system levels that are supported by IBM tape device drivers, see the IBM Spectrum Protect Supported Devices website at Supported devices for Linux.

IBM tape device drivers support only some Linux kernel levels. For information about supported kernel levels, see the Fix Central website.

IBM Spectrum Protect tape device drivers

The IBM Spectrum Protect™ server provides tape device drivers.

An IBM Spectrum Protect tape device driver is installed with the server.

AIX

You can use the generic SCSI tape device driver that is provided by the IBM® AIX® operating system to work with tape devices that are not supported by the IBM Spectrum Protect device driver. If the AIX generic SCSI tape device driver is used, the GENERICTAPE device class must be set to the device type that is specified in the DEFINE DEVCLASS command.

For the following tape devices, you can choose whether to install the IBM Spectrum Protect tape device driver or the native device driver for your operating system:

- ECART
- LTO (not from IBM)

All SCSI-attached libraries that contain tape drives from the list must use the IBM Spectrum Protect changer driver.

Tape device drivers that are acquired from other hardware vendors can be used if they are associated with the GENERICTAPE device class. Generic device drivers are not supported in write-one read-many (WORM) device classes.

Linux

You can use the IBM Spectrum Protect Passthru device driver. IBM Spectrum Protect Passthru device drivers require the Linux SCSI generic (sg) device driver along with the Linux operating system to install the kernels.

For example, you can install the IBM Spectrum Protect Passthru device driver for the following tape devices:

- ECART
- LTO (not from IBM)

All SCSI-attached libraries that contain tape drives that are not IBM labeled from the list must also use the IBM Spectrum Protect Passthru device driver.

You cannot use the generic SCSI tape (st) device driver that is provided by the Linux operating system. Therefore, the GENERICTAPE device type is not supported for the DEFINE DEVCLASS command.

Windows

You can select a Windows Hardware Qualification Lab certified native device driver instead of the IBM Spectrum Protect device driver. The Windows Hardware Qualification Lab certified native device driver can be used only for devices that have a non-IBM label and for non-IBM tape drives. For the Windows Hardware Qualification Lab certified native device driver, you can select either the IBM Spectrum Protect SCSI passthru device driver or the Windows native tape device driver. If the SCSI passthru device driver is used, the device class on the DEFINE DEVCLASS command cannot be GENERICTAPE. If the native device driver is used, the device class must be GENERICTAPE.

Special file names for tape devices

A special file name for a tape device is required for the IBM Spectrum Protect™ server to work with tape, medium changer, or removable media devices.

AIX

When a device is configured successfully, a logical file name is returned. Table 1 specifies the name of the device, also called a special file name, that corresponds to the drive or library. You can use the SMIT operating system command to get the device special file name. In the examples, *x* specifies an integer, 0 or greater.

Table 1. Device examples

Device	Device example	Logical file name
Tape drives that can be used by the IBM Spectrum Protect device driver	/dev/mtx	mtx
Tape drives that can be used by the IBM tape device driver	/dev/rmtx	rmtx
Tape drives that can be used by the IBM AIX® generic tape device driver	/dev/rmtx	rmtx
Library devices that can be used by the IBM Spectrum Protect device driver	/dev/lbx	lbx
Library devices that can be used by the IBM tape device driver	/dev/smcx	smcx

Linux

When a device is configured successfully, a logical file name is returned. Table 2 specifies the name of the device, also called the special file name, that corresponds to the drive or library. In the examples, *x* specifies an integer, 0 or greater.

Table 2. Device examples

Device	Device example	Logical file name
Tape drives that can be used by the IBM Spectrum Protect passthru device driver	/dev/sgscsi/mtx	mtx
Tape drives that can be used by the IBM lin_tape device driver	/dev/IBMtapex	IBMtapex
Library devices that can be used by the IBM Spectrum Protect passthru device driver	/dev/sgscsi/lbx	lbx
Library devices that can be used by the IBM lin_tape device driver	/dev/IBMchangerx	IBMchangerx

Windows

When a device is configured successfully, a logical file name is returned. Table 3 specifies the name of the device, also called the special file name, that corresponds to the drive or library. In the examples, *a*, *b*, *c*, *d*, and *x* specify an integer, 0 or greater, where:

- *a* specifies the target ID.
- *b* specifies the LUN.
- *c* specifies the SCSI bus ID.
- *d* specifies the port ID.

Table 3. Device examples

Device	Device example	Converted device name
Tape drives that are supported by the IBM Spectrum Protect device driver	mta.b.c.d	mta.b.c.d
Tape drives that are supported by the IBM Spectrum Protect passthru device driver	mta.b.c.d	mta.b.c.d
Tape drives that are supported by the IBM device driver	Tapex	mta.b.c.d
Library devices that are supported by the IBM Spectrum Protect device driver	lb.a.b.c.d	lba.b.c.d
Library devices that are supported by the IBM Spectrum Protect passthru device driver	lba.b.c.d	lba.b.c.d
Library devices that are supported by the IBM device driver	Changerx	lba.b.c.d

Installing and configuring tape device drivers

Before you can use tape devices with IBM Spectrum Protect™, you must install the correct tape device driver.

IBM Spectrum Protect supports all devices that are supported by IBM® tape device drivers. However, IBM Spectrum Protect does not support all the operating-system levels that are supported by IBM tape device drivers.

- Installing and configuring IBM device drivers for IBM tape devices
Install and configure an IBM tape device driver to use an IBM tape device.
- **AIX** Configuring tape device drivers on AIX systems
Review the instructions to install and configure non-IBM tape device drivers on AIX® systems.
- **Linux** Configuring tape device drivers on Linux systems
Review the following topics when you install and configure tape device drivers on Linux systems.
- **Windows** Configuring tape device drivers on Windows systems
Review the instructions to install and configure drivers for tape devices and libraries on Windows systems.

Installing and configuring IBM device drivers for IBM tape devices

Install and configure an IBM® tape device driver to use an IBM tape device.

About this task

For instructions about installing and configuring IBM tape device drivers, see the *IBM Tape Device Drivers Installation and User's Guide*.

AIX After you complete the installation procedure in the *IBM Tape Device Drivers Installation and User's Guide*, different messages are issued, depending on the device driver that you are installing. If you are installing the device driver for an IBM tape drive or library, the following messages are returned:

```
rmtx Available
```

or

```
smcx Available
```

Note the value of x, which is assigned by the IBM tape device driver. To determine the special file name of your device, issue one of the following commands:

- For tape drives, `ls -l /dev/rmt*`
- For tape libraries, `ls -l /dev/smc*`

The file name might have more characters at the end to indicate different operating characteristics, but these characters are not needed by IBM Spectrum Protect™. For IBM device drivers, use the base file name in the DEVICE parameter of the DEFINE PATH command to assign a device to a drive (/dev/rmtx) or a library (/dev/smcx).

After you install the device driver, you can use the System Management Interface Tool (SMIT) to configure non-IBM tape drives and tape libraries. Complete the following steps:

1. Run the SMIT program.

2. Click Devices.
3. Click IBM Spectrum Protect Devices.
4. Click Fibre Channel SAN Attached devices.
5. Click Discover Devices Supported by IBM Spectrum Protect. Wait for the discovery process to be completed.
6. Go back to the Fibre Channel SAN Attached devices menu, and click List Attributes of a Discovered Device.

Linux After you complete the installation procedure in the *IBM Tape Device Drivers Installation and User's Guide*, different messages are issued, depending on the device driver that you are installing. If you are installing the device driver for an IBM LTO or 3592 device, the following messages are returned:

```
IBMtapex Available
```

or

```
IBMChangerx Available
```

Note the value of x, which is assigned by the IBM tape device driver. To determine the special file name of your device, issue one of the following commands:

- For tape drives, `ls -l /dev/IBMtape*`
- For tape libraries, `ls -l /dev/IBMChange*`

The file name might have more characters at the end to indicate different operating characteristics, but these characters are not needed by IBM Spectrum Protect. For IBM device drivers, use the base file name in the DEVICE parameter of the DEFINE PATH command to assign a device to a drive (/dev/IBMtapex) or a library (/dev/IBMChangerx).

Restriction: The device type of this class must not be GENERICTAPE.

Windows For Windows operating systems, IBM Spectrum Protect provides two device drivers:

Passthru device driver

If the tape device manufacturer provides a SCSI device driver, install the IBM Spectrum Protect passthru device driver.

SCSI device driver for tape devices

If the tape device manufacturer does not provide a SCSI device driver, install the IBM Spectrum Protect SCSI device driver for tape devices. The driver file name is tsm SCSI64.sys.

For instructions about installing and configuring IBM tape device drivers, see the *IBM Tape Device Drivers Installation and User's Guide*. After you install the IBM tape device driver, the server specifies a special file name, TapeX, for IBM tape drives, or ChangerY, for IBM medium changers. For an IBM Spectrum Protect SCSI device driver or an IBM Spectrum Protect passthru device driver, you can issue the Windows operating system command, regedit, to verify the device special file name and driver. The IBM Spectrum Protect server also provides a utility to check the device for the Windows operating system. The utility, tsmdlst, is packaged with the server package. To use the utility, complete the following steps:

1. Ensure that the host bus adapter application programming interface (API) is installed.
2. To obtain device information from the host system, type:

```
tsmdlst
```

- **AIX Linux** Multipath I/O access with IBM tape devices
Multipath I/O is a technique that uses different paths to access the same physical device, for example through multiple host bus adapters (HBA) or switches. The use of the multipath technique helps to ensure that a single point of failure does not occur.

Related concepts:

Multipath I/O access with IBM tape devices

AIX

Configuring tape device drivers on AIX systems

Review the instructions to install and configure non-IBM® tape device drivers on AIX® systems.

About this task

For instructions about installing and configuring IBM tape device drivers, see the *IBM Tape Device Drivers Installation and User's Guide*.

- **AIX** SCSI and Fibre Channel devices
The IBM Spectrum Protect device definition menus and prompts in SMIT allow for the management of both SCSI and Fibre Channel (FC) attached devices.
- **AIX** Configuring IBM Spectrum Protect device drivers for autochangers
Use the following procedure to configure IBM Spectrum Protect device drivers for autochangers for non-IBM libraries.
- **AIX** Configuring IBM Spectrum Protect device drivers for tape drives
Use the following procedure to configure IBM Spectrum Protect device drivers for autochangers for vendor-acquired libraries.
- **AIX** Configuring Fibre Channel SAN-attached devices
To configure a Fibre Channel SAN-attached device, complete the procedure.

AIX

SCSI and Fibre Channel devices

The IBM Spectrum Protect™ device definition menus and prompts in SMIT allow for the management of both SCSI and Fibre Channel (FC) attached devices.

The main menu for IBM Spectrum Protect has two options:

SCSI attached devices

Use this option to configure SCSI devices that are connected to a SCSI adapter in the host.

Fibre channel system area network (SAN) attached devices

Use this option to configure devices that are connected to an FC adapter in the host. Choose one of the following attributes:

List attributes of a discovered device

Lists attributes of a device that is known to the current ODM database.

- FC Port ID:

The 24-bit FC Port ID(N(L)_Port or F(L)_Port). This is the address identifier that is unique within the associated topology where the device is connected. In the switch or fabric environments, it can be determined by the switch, with the upper 2 bytes, which are not zero. In a Private Arbitrated Loop, it is the Arbitrated Loop Physical Address(AL_PA), with the upper 2 bytes being zero. Consult with your FC vendors to find out how an AL_PA or a Port ID is assigned.

- Mapped LUN ID:

An FC to SCSI bridge (also, called a converter, router, or gateway) box. Consult with your bridge vendors about how LUNs are mapped. You should not change LUN Mapped IDs.

- WW Name:

The worldwide name of the port to which the device is attached. It is the 64-bit unique identifier that is assigned by vendors of FC components such as bridges or native FC devices. Consult with your FC vendors to find out the WWN of a port.

- Product ID:

The product ID of a device. Consult with your device vendors to determine the product ID.

Discover devices supported by IBM Spectrum Protect

This option discovers devices on an FC SAN that are supported by IBM Spectrum Protect and makes them available. If a device is added to or removed from an existing SAN environment, rediscover devices by selecting this option. Devices must be discovered first so that current values of device attributes are shown in the List Attributes of a Discovered Device option. Supported devices on FC SAN are tape drives, and autochangers. The IBM Spectrum Protect device driver ignores all other device types, such as disk.

Remove all defined devices

This option removes all FC SAN-attached IBM Spectrum Protect devices whose state is `DEFINED` in the ODM database. If necessary, rediscover devices by selecting the `Discover Devices Supported by IBM Spectrum Protect` option after the removal of all defined devices.

Remove a device

This option removes a single FC SAN-attached IBM Spectrum Protect device whose state is `DEFINED` in the ODM database. If necessary, rediscover the device by selecting the `Discover Devices Supported by IBM Spectrum Protect` option after removal of a defined device.

Configuring IBM Spectrum Protect device drivers for autochangers

Use the following procedure to configure IBM Spectrum Protect™ device drivers for autochangers for non-IBM libraries.

Procedure

Run the SMIT program to configure the device driver for each autochanger or robot:

1. Select Devices.
2. Select IBM Spectrum Protect Devices.
3. Select Library/MediumChanger.
4. Select Add a Library/MediumChanger.
5. Select the IBM Spectrum Protect-SCSI-LB for any IBM Spectrum Protect supported library.
6. Select the parent adapter to which you are connecting the device. This number is listed in the form: 00-0X, where X is the slot number location of the SCSI adapter card.
7. When prompted, enter the CONNECTION address of the device that you are installing. The connection address is a two-digit number. The first digit is the SCSI ID (the value you recorded on the worksheet). The second digit is the device's SCSI logical unit number (LUN), which is usually zero, unless otherwise noted. The SCSI ID and LUN must be separated by a comma (.). For example, a connection address of 4, 0 has a SCSI ID=4 and a LUN=0.
8. Click DO.

You receive a message (logical file name) of the form `lbX Available`. Note the value of X, which is a number that is assigned automatically by the system. Use this information to complete the Device Name field on your worksheet.

For example, if the message is `lb0 Available`, the Device Name field is `/dev/lb0` on the worksheet. Always use the `/dev/` prefix with the name provided by SMIT.

Configuring IBM Spectrum Protect device drivers for tape drives

Use the following procedure to configure IBM Spectrum Protect™ device drivers for autochangers for vendor-acquired libraries.

Procedure

Important: IBM Spectrum Protect cannot overwrite *tar* or *dd* tapes, but *tar* or *dd* can overwrite IBM Spectrum Protect tapes.
Restriction: Tape drives can be shared only when the drive is not defined or the server is not started. The MKSYB command does not work when both IBM Spectrum Protect and AIX® are sharing the same drive or drives. To use the operating system's native tape device driver with a SCSI drive, the device must be configured to AIX first and then configured to IBM Spectrum Protect. See your AIX documentation regarding these native device drivers.

Run the SMIT program to configure the device driver for each drive (including drives in libraries) as follows:

1. Select Devices.
2. Select IBM Spectrum Protect Devices.
3. Select Tape Drive.
4. Select Add a Tape Drive.
5. Select the IBM Spectrum Protect-SCSI-MT for any supported tape drive.
6. Select the adapter to which you are connecting the device. This number is listed in the form: 00-0X, where X is the slot number location of the SCSI adapter card.
7. When prompted, enter the CONNECTION address of the device you are installing. The connection address is a two-digit number. The first digit is the SCSI ID (the value you recorded on the worksheet). The second digit is the device's SCSI logical unit number (LUN), which is usually zero, unless otherwise noted. The SCSI ID and LUN must be separated by a comma (.). For example, a connection address of 4, 0 has a SCSI ID=4 and a LUN=0.
8. Click DO. You receive a message:

If you are configuring the device driver for a tape device (other than an IBM® tape drive), you receive a message (logical file name) of the form `mtX Available`. Note the value of X, which is a number that is assigned automatically by the system. Use this information to complete the Device Name field on the worksheet.

For example, if the message is `mt0 Available`, the Device Name field is `/dev/mt0` on the worksheet. Always use the `/dev/` prefix with the name provided by SMIT.

AIX

Configuring Fibre Channel SAN-attached devices

To configure a Fibre Channel SAN-attached device, complete the procedure.

Procedure

1. Run the SMIT program.
2. Select Devices.
3. Select IBM Spectrum Protect™ Devices.
4. Select Fibre Channel SAN Attached devices.
5. Select Discover Devices Supported by IBM Spectrum Protect. The discovery process can take some time.
6. Go back to the Fibre Channel menu, and select List Attributes of a Discovered Device.
7. Note the three-character device identifier, which you use when you define a path to the device to IBM Spectrum Protect. For example, if a tape drive has the identifier `mt2`, specify `/dev/mt2` as the device name.

Linux

Configuring tape device drivers on Linux systems

Review the following topics when you install and configure tape device drivers on Linux systems.

- **Linux** Configuring IBM Spectrum Protect passthru drivers for tape devices and libraries
To use the IBM Spectrum Protect Linux Passthru driver, you must complete the following steps.
- **Linux** Installing zSeries Linux Fibre Channel adapter (zfcp) device drivers
The zSeries Linux Fibre Channel adapter (zfcp) device driver is a special adapter driver on the IBM® zSeries system.
- **Linux** Information about your system's SCSI devices
Information about the devices seen by your system is available in the file `/proc/scsi/scsi`. This file contains a list of every detected SCSI device.
- **Linux** Preventing tape labels from being overwritten
The IBM Spectrum Protect Passthru device driver uses the Linux SCSI generic device driver (`sg`) to control and operate tape devices that are attached on the system. If the Linux generic SCSI tape device driver (`st`) is loaded to the kernel and configures attached tape devices, conflicts can arise over how a device is managed because the generic `sg` driver and the `st` driver can both control the same device.

Linux

Configuring IBM Spectrum Protect passthru drivers for tape devices and libraries

To use the IBM Spectrum Protect™ Linux Passthru driver, you must complete the following steps.

Procedure

1. Verify that the device is connected to your system, and is powered on and active.
2. Verify that the device is correctly detected by your system by issuing this command:

```
cat /proc/scsi/scsi
```

3. Ensure that both the IBM Spectrum Protect device driver package (`tmsmcsi`) and the storage server package are installed.
4. There are two driver configuration methods available in the IBM Spectrum Protect device driver package: `autoconf` and `tmsmcsi`. Both of these methods complete the following tasks:
 - Load the Linux SCSI generic driver (`sg`) to the kernel.
 - Create necessary special files for the Passthru driver.
 - Create device information files for tape devices (`/dev/tmsmcsi/mtinfo`) and libraries (`/dev/tmsmcsi/lbinfo`).
5. Run the configuration method that you prefer (`autoconf` or `tmsmcsi`) for the IBM Spectrum Protect Passthru driver.
 - To run the `autoconf` configuration method, issue the following command:

```
autoconf
```

- o To run the tsm SCSI configuration method, complete the following steps:
 - a. Copy the two sample configuration files that are in the installation directory from *mt.conf.smp* and *lb.conf.smp* to *mt.conf* and *lb.conf*, respectively.
 - b. Edit the *mt.conf* and *lb.conf* files. Add one stanza (as shown in the example at the start of the file) for each SCSI target, ID, and LUN combination. Each combination of SCSI target, ID, and LUN entries correspond to a tape drive or library you want configured. Make sure that the files meet these requirements:
 - Remove the example that is at the start of the files.
 - There must be a new line between each stanza.
 - There must be one new line after the last stanza.
 - Ensure that there are no number signs (#) in either file.
 - c. Run the *tsmscsi* script from the device driver installation directory.
- 6. Verify that the device is configured properly by viewing the text files for tape devices (*/dev/tsmscsi/mtinfo*) and libraries (*/dev/tsmscsi/lbinfo*).
- 7. Determine the special file names for the tape drives and libraries:

- o To determine the names for tape devices, issue the following command:

```
> ls /dev/tsmscsi/mt*
```

- o To determine the names for libraries, issue the following command:

```
> ls /dev/tsmscsi/lb*
```

This information helps you identify which of the */dev/tsmscsi/mtx* and */dev/tsmscsi/lbx* special file names to provide the server when you issue a *DEFINE PATH* command.

What to do next

If you restart the host system, you must rerun the *autoconf* or *tsmscsi* script to reconfigure IBM Spectrum Protect devices. If you restart the IBM Spectrum Protect server instance, you do not have to reconfigure devices. In general, the Linux SCSI generic driver is preinstalled to the kernel. To verify that the driver is in the kernel, issue the following command:

```
> lsmod | grep sg
```

If the driver is not in the kernel, issue the *modprobe sg* command to load the *sg* driver into the kernel.

Linux

Installing zSeries Linux Fibre Channel adapter (zfcp) device drivers

The zSeries Linux Fibre Channel adapter (zfcp) device driver is a special adapter driver on the IBM® zSeries system.

About this task

IBM Spectrum Protect™ and IBM tape device drivers can run on zSeries platforms with Linux operating systems in 64-bit environments, and support most original equipment manufacturer (OEM) and IBM tape devices with Fibre Channel interfaces.

For more information about the *zfcp* driver, see the IBM Redpaper™, *Getting Started with zSeries Fibre Channel Protocol*, which is available at IBM Redbooks®.

Procedure

1. Load the *qdio* module.
2. Install the *zfcp* driver.
3. Map the Fibre Channel Protocol (FCP) and configure the *zfcp* driver.
4. Install and configure the IBM tape device driver.

Linux

Information about your system's SCSI devices

Information about the devices seen by your system is available in the file */proc/scsi/scsi*. This file contains a list of every detected SCSI device.

The following device information is available: the host number, channel number, SCSI ID, Logical Unit number, vendor, firmware level, type of device, and the SCSI mode. For example, if a system contains some StorageTek and IBM® libraries, a SAN Gateway, and some Quantum DLT drives, the `/proc/scsi/scsi` file will look similar to this:

```
Attached devices:
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: STK      Model: 9738      Rev: 2003
  Type:  Medium Changer      ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: PATHLIGHT Model: SAN Gateway  Rev: 32aC
  Type: Unknown      ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: QUANTUM Model: DLT7000      Rev: 2560
  Type: Sequential-Access  ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 01 Lun: 04
  Vendor: IBM      Model: 7337      Rev: 1.63
  Type:  Medium Changer      ANSI SCSI revision: 02
```

Linux

Preventing tape labels from being overwritten

The IBM Spectrum Protect™ Passthru device driver uses the Linux SCSI generic device driver (`sg`) to control and operate tape devices that are attached on the system. If the Linux generic SCSI tape device driver (`st`) is loaded to the kernel and configures attached tape devices, conflicts can arise over how a device is managed because the generic `sg` driver and the `st` driver can both control the same device.

About this task

If the `st` driver controls devices that are used by IBM Spectrum Protect, IBM Spectrum Protect internal tape labels can be overwritten and data can be lost. If an application uses the `st` driver to control devices and the non-rewind option is not specified, tapes are automatically rewound following completion of an operation. The auto-rewind operation relocates the tape header position to the beginning of the tape. If the tape remains loaded in the drive, the next non-IBM Spectrum Protect write operation overwrites the IBM Spectrum Protect tape label because the label is at the beginning of the tape.

To prevent IBM Spectrum Protect labels from being overwritten, which can result in data loss, ensure that only the IBM Spectrum Protect Passthru driver controls devices that are used by IBM Spectrum Protect. Remove the `st` driver from the kernel or, if the driver is used by some applications on the system, delete the special files that correspond to IBM Spectrum Protect devices so that the `st` driver can no longer control them.

If you are using the IBM tape device driver to control devices on your system, you might encounter the same issues with device driver control conflicts. Review your IBM tape documentation to determine how to resolve this issue and prevent data loss.

Remove the `st` driver

If no other applications on the system use `st` devices, remove the `st` driver from the kernel. Issue the following command to unload the `st` driver:

```
rmmod st
```

Delete device special files that correspond to IBM Spectrum Protect devices

If there are applications that require use of the `st` driver, delete the special files that correspond to IBM Spectrum Protect devices. These special files are generated by the `st` driver. When they are eliminated, the `st` driver can no longer control the corresponding IBM Spectrum Protect devices. Device special file names for tape drives appear in the `/dev/` directory. Their names have the form `/dev/[n]st[0-1024][l][m][a]`.

List the `st` drive special file names and IBM Spectrum Protect device special file names by using the `ls` command. Based on the output of the device sequences, you can find devices in the `st` devices list matching those in the IBM Spectrum Protect devices list. The `rm` command can then be used to delete `st` devices.

Issue the following commands to list the `st` and IBM Spectrum Protect devices:

```
ls -l /dev/*st*
ls -l /dev/tmsmcsi/mt*
```

Delete the `st` devices with the `rm` command:

```
rm /dev/*st*
```


Configuring tape device drivers on Windows systems

Review the instructions to install and configure drivers for tape devices and libraries on Windows systems.

- **Windows** Preparing to use the IBM Spectrum Protect passthru driver for tape devices and libraries
To use the IBM Spectrum Protect Windows passthru device driver for tape devices and libraries, you must install the driver and obtain the device names for the server to use.
- **Windows** Configuring the IBM Spectrum Protect SCSI driver for tape devices and libraries
If the manufacturer of a tape drive or tape library does not provide a SCSI device driver, you must install the IBM Spectrum Protect SCSI device driver.

Preparing to use the IBM Spectrum Protect passthru driver for tape devices and libraries

To use the IBM Spectrum Protect™ Windows passthru device driver for tape devices and libraries, you must install the driver and obtain the device names for the server to use.

Before you begin

1. Determine whether the manufacturer of the tape device or tape library provides a device driver.
2. If the manufacturer provides a device driver package, download the package and install it.
3. Configure the SCSI device driver by following the manufacturer's instructions.

Procedure

1. Install the IBM Spectrum Protect passthru device driver.
2. Obtain the device names that the server must use by taking one of the following actions:
 - On the server, run the QUERY SAN command. The output shows all devices names and their associated device serial numbers.
 - In the server directory, run the tsmdlst.exe utility. The output shows all devices names, their associated serial numbers, and associated device locations.
 - At the Windows system command prompt, run the regedit command. From the output, obtain the device file names based on the device locations. The location consists of the port ID, SCSI bus ID, LUN ID, and SCSI target ID. The IBM Spectrum Protect device file name has a format of mtA.B.C.C for tape drives and lbA.B.C.D for tape libraries, where:
 - A is the SCSI target ID.
 - B is the LUN ID.
 - C is the SCSI bus ID.
 - D is the port ID.

Configuring the IBM Spectrum Protect SCSI driver for tape devices and libraries

If the manufacturer of a tape drive or tape library does not provide a SCSI device driver, you must install the IBM Spectrum Protect™ SCSI device driver.

About this task

The IBM Spectrum Protect SCSI device driver file name is tsm SCSI64.sys.

Procedure

1. Locate the device in the Device Manager console (devmgmt.msc) and select it. Tape drives are listed under Tape drives, and medium changers are under Medium Changers.

2. Configure the device for use by the tsm SCSI64.sys device driver:
 - a. Right-click the device and click Update Driver Software.
 - b. Click Browse my computer for driver software.
3. Click Let me pick from a list of device drivers on my computer.
4. Click Next.
5. Select the appropriate option:
 - a. For a tape drive, select IBM Spectrum Protect for Tape Drives.
 - b. For a medium changer, select IBM Spectrum Protect for Medium Changers.
6. Click Next.
7. Click Close.
8. Verify that the device was configured correctly for the tsm SCSI64 device driver:
 - a. Right-click on the device and click Properties.
 - b. Click the Driver tab and Driver Details. The Driver Details window shows the device driver that is controlling the device.

Configuring libraries for use by a server

To use a library or libraries for storage for an IBM Spectrum Protect™ server, you must first set up the devices on the server system.

Before you begin

1. Attach devices to the server hardware. Follow the instructions in [Attaching an automated library device to your system](#).
2. Select the tape device drivers. Follow the instructions in [Selecting a tape device driver](#).
3. Install and configure the tape device drivers. Follow the instructions in [Installing and configuring tape device drivers](#).
4. Determine the device names that are needed to define the library to the server. Follow the instructions in [Special file names for tape devices](#).

Procedure

1. Define the library and the path from the server to the library. Follow the instructions in [Defining libraries](#).
2. Define the drives in the library. Follow the instructions in [Defining drives](#).

For SCSI libraries, you can use the `PERFORM LIBACTION` command to define drives and paths for a library in one step, instead of completing both steps 2 and 3. To use the `PERFORM LIBACTION` command to define drives and paths for a library, the `SANDISCOVERY` option must be supported and enabled.

3. Define a path from the server to each drive by using the `DEFINE PATH` command.
4. Define a device class. Follow the instructions in [Defining tape device classes](#).

Device classes specify the recording formats for drives and classify them according to type. Use the default value, `FORMAT=DRIVE` as the recording format only if all the drives that are associated with the device class can read and write to all of the media.

For example, you have a mix of Ultrium Generation 3 and Ultrium Generation 4 drives, but you have only Ultrium Generation 3 media. You can specify `FORMAT=DRIVE` because both the Generation 4 and Generation 3 drives can read from and write to Generation 3 media.

5. Define a storage pool by using the `DEFINE STGPOOL` command.

Consider the following key choices for defining storage pools:

- o Scratch volumes are empty volumes that are available for use. If you specify a value for the maximum number of scratch volumes in the storage pool, the server can choose from the scratch volumes available in the library.

If you do not allow scratch volumes, you must complete the extra step of explicitly defining each volume to be used in the storage pool. Also, specify the `MAXSCRATCH=0` parameter when you define the storage pool so that scratch volumes are not used.

- o The default setting for primary storage pools is collocation by group. The default for copy storage pools and active-data pools is disablement of collocation. The server uses *collocation* to keep all files that belong to a group of client nodes, a single client node, a client file space, or a group of client file spaces on a minimal number of volumes. If collocation is disabled for a storage pool and clients begin storing data, you cannot easily change the data in the pool so that it is collocated.

6. Check in and label library volumes. Follow the instructions in Checking volumes into an automated library and Labeling tape volumes.

Ensure that enough volumes in the library are available to the server. Keep enough labeled volumes on hand so that you do not run out during an operation such as client backup. Label extra scratch volumes for any potential recovery operations that you might have later.

The procedures for checking in and labeling volumes are the same whether the library contains drives of a single device type, or drives of multiple device types. You can use the CHECKIN LIBVOLUME command to check in volumes that are already labeled. Or, if you want to label and check in volumes with one step, issue the LABEL LIBVOLUME command.

Libraries with multiple device types: If your library has drives of multiple device types, and you defined two libraries to the IBM Spectrum Protect server, the two defined libraries represent one physical library. You must check in tape volumes separately to each defined library. Ensure that you check in volumes to the correct IBM Spectrum Protect library.

What to do next

Verify your device definitions to ensure that everything is configured correctly. Use a QUERY command to review information about each storage object.

When you review the results of the QUERY DRIVE command, verify that the device type for the drive is what you expect. If a path is not defined, the drive device type is listed as UNKNOWN and if the wrong path is used, GENERIC_TAPE or another device type is shown. This step is especially important when you are using mixed media.

Optionally, configure library sharing. Follow the instructions in Configuring library sharing.

- Defining tape devices
Before you can back up or migrate data to tape, you must define a tape device to the IBM Spectrum Protect.
- Configuring library sharing
Multiple IBM Spectrum Protect servers can share storage devices by using a storage area network (SAN). You set up one server as the library manager and the other servers as library clients.

Related reference:

- [CHECKIN LIBVOLUME](#) (Check a storage volume into a library)
- [LABEL LIBVOLUME](#) (Label a library volume)
- [PERFORM LIBACTION](#) (Define or delete all drives and paths for a library)

Related information:

- [DEFINE STGPOOL](#) (Define a volume in a storage pool)

Defining tape devices

Before you can back up or migrate data to tape, you must define a tape device to the IBM Spectrum Protect™.

- Defining libraries and drives
A tape library can include one or more tape drives. Learn how to define libraries, drives, and paths to the IBM Spectrum Protect server.
- Defining tape device classes
A device class defines a set of characteristics that are used by a set of volumes that can be created in a storage pool. You must define a device class for a tape device to ensure that the server can use the device.

Defining libraries and drives

A tape library can include one or more tape drives. Learn how to define libraries, drives, and paths to the IBM Spectrum Protect™ server.

- Defining libraries
Before you can use a drive, you must define the library to which the drive belongs.
- Defining drives
To inform the server about a drive that can be used to access storage volumes, issue the DEFINE DRIVE command, followed by the DEFINE PATH command.

Defining libraries

Before you can use a drive, you must define the library to which the drive belongs.

Procedure

1. Define the library by using the DEFINE LIBRARY command.

For example, if you have an IBM TS3500 tape library, you can define a library that is named ROBOTMOUNT by using the following command:

```
define library robotmount libtype=scsi
```

If you require library sharing or LAN-free data movement, see the following information:

- o Configuring library sharing
- o Configuring LAN-free data movement

2. Define a path from the server to the library by using the DEFINE PATH command. When you specify the DEVICE parameter, enter the device special file name. This name is required by the server to communicate with tape drives, medium changer, and removable media devices. For more information about device special file names, see [Special file names for tape devices](#).

```
define path server1 robotmount srctype=server desttype=library  
device=/dev/lb0
```

Linux

```
define path server1 robotmount srctype=server desttype=library  
device=/dev/tmscsi/lb0
```

Windows

```
define path server1 robotmount srctype=server desttype=library  
device=lb0.0.1.0
```

- Defining SCSI libraries on a SAN

For a library type of SCSI on a SAN, the server can track the library's serial number. With the serial number, the server can confirm the identity of the device when you define the path or when the server uses the device.

Related information:

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

Defining drives

To inform the server about a drive that can be used to access storage volumes, issue the DEFINE DRIVE command, followed by the DEFINE PATH command.

Before you begin

A *drive object* represents a drive mechanism within a library that uses removable media. For devices with multiple drives, including automated libraries, you must define each drive separately and associate it with a library. Drive definitions can include such information as the element address for drives in SCSI, how often a tape drive is cleaned, and whether the drive is online.

IBM Spectrum Protect™ supports tape drives that can be stand-alone or that can be part of an automated library. The preferred method is to configure the tape solution by using automated libraries.

About this task

When you issue the DEFINE DRIVE command, you must provide some or all of the following information:

Library name

The name of the library in which the drive is located.

Drive name

The name that is assigned to the drive.

Serial number

The serial number of the drive. The serial number parameter applies only to drives in SCSI. With the serial number, the server can confirm the identity of the device when you define the path or when the server uses the device.

You can specify the serial number if you choose. The default is to enable the server to obtain the serial number from the drive itself at the time that the path is defined. If you specify the serial number, the server confirms that the serial number is correct when you define the path to the drive. When you define the path, you can set the AUTODETECT=YES parameter to enable the server to correct the serial number if the number that it detects does not match what you entered when you defined the drive. As a best practice, specify the AUTODETECT=YES parameter to automatically update the serial number for the drive in the database when the path is defined.

Depending on the capabilities of the drive, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device. See Impacts of device changes on the SAN.

Element address

The element address of the drive. The ELEMENT parameter applies only to drives in SCSI libraries. The element address is a number that indicates the physical location of a drive within an automated library. The server needs the element address to connect the physical location of the drive to the drive's SCSI address. The server can obtain the element address from the drive when you define the path, or you can specify the element number when you define the drive. As a best practice, specify the ELEMENT=AUTODETECT parameter for the server to automatically detect the element number when the path to the drive is defined.

Depending on the capabilities of the library, the server might not be able to automatically detect the element address. In this case, you must supply the element address when you define the drive, if the library has more than one drive. To obtain the element address, go to the IBM® Support Portal for IBM Spectrum Protect.

Tip: IBM tape device drivers and non-IBM tape device drivers generate different device files and formats:

- For IBM, device names begin with rmt followed by an integer, for example, /dev/rmt0.
- For IBM Spectrum Protect tape device drivers, tape device names begin with mt followed by an integer, for example /dev/mt0.

You must use the correct device file when you define a path.

Procedure

1. Assign a drive to a library by issuing the DEFINE DRIVE command.
2. To make the drive usable by the server, issue the DEFINE PATH command.

For examples about configuring libraries, paths, and drives, see Example: Configure a SCSI or virtual tape library with a single drive device type and Example: Configure a SCSI or virtual tape library with multiple drive device types.

Defining tape device classes

A device class defines a set of characteristics that are used by a set of volumes that can be created in a storage pool. You must define a device class for a tape device to ensure that the server can use the device.

Before you begin

You must define libraries and drives to the server before you define device classes.

About this task

For a list of supported devices and valid device class formats, see the IBM Spectrum Protect™ Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

You can define multiple device classes for each device type. For example, you might want to specify different attributes for different storage pools that use the same type of tape drive. Variations might be required that are not specific to the device, but rather to how you want to use the device (for example, mount retention or mount limit).

Guidelines:

- One device class can be associated with multiple storage pools, but each storage pool is associated with only one device class.
- SCSI libraries can include tape drives of more than one device type. When you define the device class in this environment, you must declare a value for the FORMAT parameter.

For more information, see Mixed device types in libraries.

Procedure

To define a device class, use the DEFINE DEVCLASS command with the DEVTYPE parameter, which assigns a device type to the device class.

Results

If you include the DEVCONFIG option in the dsmserv.opt file, the files that you specify with that option are automatically updated with the results of the DEFINE DEVCLASS, UPDATE DEVCLASS, and DELETE DEVCLASS commands.

- Defining LTO device classes
To prevent problems when you mix different generations of LTO drives and media in a single library, review the restrictions. Also, review the restrictions for LTO drive encryption.
- Defining 3592 device classes
Device class definitions for 3592, TS1130, TS1140, TS1150, and later devices include parameters for faster volume-access speeds and drive encryption. To prevent problems when mixing different generations of 3592 and TS1130 and later drives in a library, review the guidelines.

Related reference:

[DEFINE DEVCLASS \(Define a device class\)](#)

Related information:

[QUERY DEVCLASS \(Display information on one or more device classes\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Defining LTO device classes

To prevent problems when you mix different generations of LTO drives and media in a single library, review the restrictions. Also, review the restrictions for LTO drive encryption.

- Mixing LTO drives and media in a library
When you mix different generations of LTO drives and media, you must consider the read/write capabilities of each generation. The preferred method is to configure a different device class for each generation of media.
- Mount limits in LTO mixed-media environments
In a mixed-media library, in which multiple device classes point to the same library, compatible drives are shared between storage pools. Ensure that you set an appropriate value for the MOUNTLIMIT parameter in each of the device classes.
- Enabling and disabling drive encryption for LTO Generation 4 or later tape drives
IBM Spectrum Protect™ supports the three types of drive encryption that are available with LTO Generation 4 or later drives: Application, System, and Library. These methods are defined through the hardware.

Mixing LTO drives and media in a library

When you mix different generations of LTO drives and media, you must consider the read/write capabilities of each generation. The preferred method is to configure a different device class for each generation of media.

About this task

If you are considering mixing different generations of LTO media and drives, review the following restrictions:

Table 1. Read/write capabilities for different generations of LTO drives

Drives	Generatio n 1 media	Generatio n 2 media	Generatio n 3 media	Generatio n 4 media	Generatio n 5 media	Generatio n 6 media	Generatio n 7 media	Generatio n M8 media	Generatio n 8 media
--------	------------------------	------------------------	------------------------	------------------------	------------------------	------------------------	------------------------	----------------------------	------------------------

Drives	Generatio n 1 media	Generatio n 2 media	Generatio n 3 media	Generatio n 4 media	Generatio n 5 media	Generatio n 6 media	Generatio n 7 media	Generatio n M8 media	Generatio n 8 media
Generation 1	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Generation 2	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Generation 3	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a
Generation 4	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a
Generation 5	n/a	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a
Generation 6	n/a	n/a	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a
Generation 7	n/a	n/a	n/a	n/a	Read access	Read/write access	Read/write access	n/a	n/a
Generation 8	n/a	n/a	n/a	n/a	n/a	n/a	Read/write access	Read/write access	Read/write access

Example

If you are mixing different types of drives and media, configure different device classes: one for each type of media. To specify the media type, use the `FORMAT` parameter in each of the device class definitions. (Do not specify `FORMAT=DRIVE`.) For example, if you are mixing Ultrium Generation 5 and Ultrium Generation 6 drives, specify `FORMAT=ULTRIUM5C` (or `ULTRIUM5`) for the Ultrium Generation 5 device class, and `FORMAT=ULTRIUM6C` (or `ULTRIUM6`) for the Ultrium Generation 6 device class.

In this example, both device classes can point to the same library with Ultrium Generation 5 and Ultrium Generation 6 drives. The drives are shared between the two storage pools. One storage pool uses the first device class and Ultrium Generation 5 media exclusively. The other storage pool uses the second device class and Ultrium Generation 6 media exclusively. Because the two storage pools share a single library, Ultrium Generation 5 media can be mounted on Ultrium Generation 6 drives as they become available during mount point processing.

If you mix older read-only media generations with newer read/write media in a single library, you must mark the read-only media as read-only and check out all read-only scratch media. For example, if you are mixing Ultrium Generation 4 with Ultrium Generation 6 drives and media in a single library, you must mark the Generation 4 media as read-only. In addition, you must check out all Generation 4 scratch volumes.

Mount limits in LTO mixed-media environments

In a mixed-media library, in which multiple device classes point to the same library, compatible drives are shared between storage pools. Ensure that you set an appropriate value for the `MOUNTLIMIT` parameter in each of the device classes.

For example, in a mixed media library that contains Ultrium Generation 1 and Ultrium Generation 2 drives and media, Ultrium Generation 1 media can be mounted in Ultrium Generation 2 drives.

Consider the example of a mixed library that consists of the following drives and media:

- Four LTO Ultrium Generation 1 drives and LTO Ultrium Generation 1 media
- Four LTO Ultrium Generation 2 drives and LTO Ultrium Generation 2 media

You created the following device classes:

- LTO Ultrium Generation 1 device class `LTO1CLASS` specifying `FORMAT=ULTRIUM1C`
- LTO Ultrium Generation 2 device class `LTO2CLASS` specifying `FORMAT=ULTRIUM2C`

You also created the following storage pools:

- LTO Ultrium Generation 1 storage pool `LTO1POOL` based on device class `LTO1CLASS`

- LTO Ultrium Generation 2 storage pool LTO2POOL based on device class LTO2CLASS

The number of mount points available for use by each storage pool is specified in the device class by using the MOUNTLIMIT parameter. The MOUNTLIMIT parameter in the LTO2CLASS device class must be set to 4 to match the number of available drives that can mount only LTO7 media. The MOUNTLIMIT parameter in the LTO1CLASS device class must be set to a value that is greater than the number of available drives (5 or possibly 6) to adjust for the fact that Ultrium Generation 1 media can be mounted in Ultrium Generation 7 drives. The optimal value for MOUNTLIMIT depends on workload and storage pool access patterns.

Monitor and adjust the MOUNTLIMIT setting to suit changing workloads. If the MOUNTLIMIT for LTO1POOL is set too high, mount requests for the LTO2POOL might be delayed or fail because the Ultrium Generation 2 drives are used to satisfy Ultrium Generation 1 mount requests. In the worst scenario, too much competition for Ultrium Generation 2 drives might cause mounts for Generation 2 media to fail with the following message:

```
ANR8447E No drives are currently available in the library.
```

If the MOUNTLIMIT value for LTO1POOL is not set high enough, mount requests that might be satisfied by LTO Ultrium Generation 2 drives are delayed.

Restriction: Restrictions apply when you mix Ultrium Generation 1 with Ultrium Generation 2 or Generation 3 drives because of how mount points are allocated. For example, processes that require multiple mount points that include both Ultrium Generation 1 and Ultrium Generation 2 volumes might try to reserve Ultrium Generation 2 drives only, even when one mount can be satisfied by an available Ultrium Generation 6 drive. Processes that behave in this manner include the MOVE DATA and BACKUP STGPOOL commands. These processes wait until the required number of mount points can be satisfied with Ultrium Generation 2 drives.

Related reference:

- [BACKUP STGPOOL](#) (Back up primary storage pool data to a copy storage pool)
- [DEFINE DEVCLASS](#) (Define a device class)
- [MOVE DATA](#) (Move files on a storage pool volume)

Enabling and disabling drive encryption for LTO Generation 4 or later tape drives

IBM Spectrum Protect™ supports the three types of drive encryption that are available with LTO Generation 4 or later drives: Application, System, and Library. These methods are defined through the hardware.

About this task

The DRIVEENCRYPTION parameter on the DEFINE DEVCLASS command specifies whether drive encryption is allowed for IBM and HP LTO Generation 4 or later, Ultrium 4, and Ultrium 4C formats. This parameter ensures IBM Spectrum Protect compatibility with hardware encryption settings for empty volumes. You cannot use this parameter for storage pool volumes that are full or are filling.

IBM Spectrum Protect supports the Application method of encryption with IBM and HP LTO-4 or later drives. Only IBM LTO-4 or later supports the System and Library methods. The Library method of encryption can be used only if your system hardware (for example, IBM TS3500) supports it.

Restriction: You cannot use drive encryption with write-once, read-many (WORM) media.

The Application method is defined through the hardware. To use the Application method, in which IBM Spectrum Protect generates and manages encryption keys, set the DRIVEENCRYPTION parameter to ON. This action enables data encryption for empty volumes. If the parameter is set to ON and the hardware is configured for another encryption method, backup operations fail.

Procedure

The following simplified example shows the steps that you would take to enable and disable data encryption for empty volumes in a storage pool:

1. Define a library by issuing the DEFINE LIBRARY command:

```
define library 3584 libtype=SCSI
```

2. Define a device class, LTO_ENCRYPT, by issuing the DEFINE DEVCLASS command and specifying IBM Spectrum Protect as the key manager:


```
define devclass lto_encrypt library=3584 devtype=lto driveencryption=on
```

3. Define a storage pool by issuing the DEFINE STGPOOL command:

```
define stgpool lto_encrypt_pool lto_encrypt
```

4. To disable encryption on new volumes, set the DRIVEENCRYPTION parameter to OFF. The default value is ALLOW. Drive encryption for empty volumes is allowed if another method of encryption is enabled.

Related concepts:

Tape encryption methods

Defining 3592 device classes

Device class definitions for 3592, TS1130, TS1140, TS1150, and later devices include parameters for faster volume-access speeds and drive encryption. To prevent problems when mixing different generations of 3592 and TS1130 and later drives in a library, review the guidelines.

- **Mixing generations of 3592 drives and media in a single library**
For optimal performance, do not mix generations of 3592 media in a single library. Media problems can result when different drive generations are mixed. For example, IBM Spectrum Protect™ might not be able to read a volume's label.
- **Controlling data-access speeds for 3592 volumes**
You can optimize the storage capacity and improve data-access speeds when you create volumes. By partitioning data into storage pools that have volumes, you can specify the scale capacity percentage to provide maximum storage capacity, or to provide fast access to the volume.
- **Enabling and disabling 3592 Generation 2 and later drive encryption**
With IBM Spectrum Protect, you can use the following types of drive encryption with drives that are 3592 Generation 2 and later: Application, System, and Library. These methods are defined through the hardware.

Mixing generations of 3592 drives and media in a single library

For optimal performance, do not mix generations of 3592 media in a single library. Media problems can result when different drive generations are mixed. For example, IBM Spectrum Protect™ might not be able to read a volume's label.

About this task

The following table shows read/write interoperability for drive generations.

Drives	Generation 1 format	Generation 2 format	Generation 3 format	Generation 4 format	Generation 5 format
Generation 1	Read/write access	n/a	n/a	n/a	n/a
Generation 2	Read/write access	Read/write access	n/a	n/a	n/a
Generation 3	Read-only access	Read/write access	Read/write access	n/a	n/a
Generation 4	n/a	Read only	Read/write access	Read/write access	n/a
Generation 5	n/a	n/a	Read access	Read/write access	Read/write access

If you must mix generations of drives in a library, review the example and restrictions to help prevent problems.

Table 1. Mixing generations of drives

Library type	Example and restrictions
--------------	--------------------------

Library type	Example and restrictions
SCSI	<p>Define a new storage pool and device class for the latest drive generation. For example, suppose that you have a storage pool and device class for 3592-2. The storage pool contains all the media that were written in Generation 2 format. Suppose that the value of the FORMAT parameter in the device class definition is set to 3952-2 (not DRIVE). You add Generation 3 drives to the library. Complete the following steps:</p> <ol style="list-style-type: none"> 1. In the new device-class definition for the Generation 3 drives, set the value of the FORMAT parameter to 3592-3 or 3592-3C. Do not specify DRIVE. 2. In the definition of the storage pool that is associated with Generation 2 drives, update the MAXSCRATCH parameter to 0, for example: <pre data-bbox="542 478 1024 506">update stgpool genpool2 maxscratch=0</pre> <p>This method allows both generations to use their optimal format and minimizes potential media problems that can result from mixing generations. However, it does not resolve all media issues. For example, competition for mount points and mount failures might result. (To learn more about mount point competition in the context of 3592 drives and media, see Defining 3592 device classes.)</p> <p>Restriction: The following list describes media restrictions:</p> <ul style="list-style-type: none"> • CHECKIN LIBVOL: The issue is using the CHECKLABEL=YES option. If the label is written in a Generation 3 or later format, and you specify the CHECKLABEL=YES option, drives of previous generations fail by using this command. To avoid the issue, specify CHECKLABEL=BARCODE. • LABEL LIBVOL: When the server tries to use drives of a previous generation to read the label that is written in a Generation 3 or later format, the LABEL LIBVOL command fails unless OVERWRITE=YES is specified. Verify that the media that is being labeled with OVERWRITE=YES does not have any active data. • CHECKOUT LIBVOL: When IBM Spectrum Protect verifies the label (CHECKLABEL=YES) as a Generation 3 or later format, and read drives of previous generations, the command fails. To avoid this issue, specify CHECKLABEL=NO.

Related reference:

- [CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)
- [CHECKOUT LIBVOLUME \(Check a storage volume out of a library\)](#)
- [LABEL LIBVOLUME \(Label a library volume\)](#)
- [UPDATE STGPOOL \(Update a storage pool\)](#)

Controlling data-access speeds for 3592 volumes

You can optimize the storage capacity and improve data-access speeds when you create volumes. By partitioning data into storage pools that have volumes, you can specify the scale capacity percentage to provide maximum storage capacity, or to provide fast access to the volume.

About this task

To reduce media capacity, specify the SCALECAPACITY parameter when you define the device class by using the DEFINE DEVCLASS command or when you update the device class by using the UPDATE DEVCLASS command.

Specify a percentage value of 20, 90, or 100. A value of 20 percent provides the fastest access time, and 100 percent provides the largest storage capacity. For example, if you specify a scale capacity of 20 for a 3592 device class without compression, a 3592 volume in that device class would store 20 percent of its full capacity of 300 GB, or about 60 GB.

Scale capacity takes effect only when data is first written to a volume. Updates to the device class for scale capacity do not affect volumes that already have data written to them until the volume is returned to scratch status.

Related reference:

- [DEFINE DEVCLASS \(Define a device class\)](#)

Related information:

- [UPDATE DEVCLASS \(Update a device class\)](#)

Enabling and disabling 3592 Generation 2 and later drive encryption

With IBM Spectrum Protect™, you can use the following types of drive encryption with drives that are 3592 Generation 2 and later: Application, System, and Library. These methods are defined through the hardware.

About this task

The DRIVEENCRYPTION parameter on the DEFINE DEVCLASS command specifies whether drive encryption is allowed for drives that are 3592 Generation 2 and later. Use this parameter to ensure IBM Spectrum Protect compatibility with hardware encryption settings for empty volumes. You cannot use this parameter for storage pool volumes that are full or are filling.

- To use the Application method, in which IBM Spectrum Protect generates and manages encryption keys, set the DRIVEENCRYPTION parameter to ON. This enables the encryption of data for empty volumes. If the parameter is set to ON and if the hardware is configured for another encryption method, backup operations fail.
- To use the Library or System methods of encryption, set the parameter to ALLOW. This specifies that IBM Spectrum Protect is not the key manager for drive encryption, but allows the hardware to encrypt the volume's data through one of the other methods. Specifying this parameter does not automatically encrypt volumes. Data can be encrypted only by specifying the ALLOW parameter and configuring the hardware to use one of these methods.

The DRIVEENCRYPTION parameter is optional. The default value is to allow the Library or System methods of encryption.

Procedure

The following simplified example shows how to encrypt data for empty volumes in a storage pool, by using IBM Spectrum Protect as the key manager:

1. Define a library by issuing the DEFINE LIBRARY command. For example, issue the following command:

```
define library 3584 libtype=SCSI
```

2. Define a device class, 3592_ENCRYPT, by issuing the DEFINE DEVCLASS command and specifying the value ON for the DRIVEENCRYPTION parameter. For example, issue the following command:

```
define devclass 3592_encrypt library=3584 devtype=3592 driveencryption=on
```

3. Define a storage pool. For example, issue the following command:

```
define stgpool 3592_encrypt_pool 3592_encrypt
```

What to do next

To disable any method of encryption on new volumes, set the DRIVEENCRYPTION parameter to OFF. If the hardware is configured to encrypt data through either the Library or System method and DRIVEENCRYPTION is set to OFF, backup operations fail.

Configuring library sharing

Multiple IBM Spectrum Protect™ servers can share storage devices by using a storage area network (SAN). You set up one server as the library manager and the other servers as library clients.

Before you begin

Ensure that your systems meet licensing requirements for library sharing. An entitlement for IBM Spectrum Protect for SAN is required for each IBM Spectrum Protect server that is configured as a library client or a library manager in a SAN environment.

About this task

With LAN-free data movement, IBM Spectrum Protect client systems can directly access storage devices that are defined to an IBM Spectrum Protect server. Storage agents are installed and configured on the client systems to perform the data movement.

To set up library sharing, you must define one IBM Spectrum Protect server as the library manager for your shared library configuration. Then, you define other IBM Spectrum Protect servers as library clients that communicate and request storage

resources from the library manager. The library manager server must be at the same version or a later version as the server or servers that are defined as library clients.

Procedure

To complete the following steps to share library resources on a SAN among IBM Spectrum Protect servers, complete the following steps:

1. Set up server-to-server communications.

To share a storage device on a SAN, define servers to each other by using the cross-define function. Each server must have a unique name.

2. Define a shared library and set up tape devices on the server systems.

Use the procedure that is described in [Configuring libraries for use by a server](#) to define a library for use in the shared environment. Modify the procedure to define the library as shared, by specifying the SHARED=YES parameter for the DEFINE LIBRARY command.

3. Define the library manager server.
4. Define the shared library on the server that is the library client.
5. From the library manager server, define paths from the library client to each drive that the library client can access. The device name must reflect the way that the library client system recognizes the tape device. A path from the library manager to each tape drive must be defined in order for the library client to use the drive.

To avoid problems, ensure that all drive path definitions that are defined for the library manager are also defined for each library client.


For example, if the library manager defines three tape drives, the library client must also define three tape drives. To limit the number of tape drives that a library client can use at a time, use the MOUNTLIMIT parameter of the device class on the library client.

6. Define device classes for the shared library.

The preferred method is to make the device class names the same on both servers to avoid confusion when you define multiple device classes with the same device type and library parameters. Some operations, such as database backup, use the device class name to identify the data for backup.

The device class parameters that are specified on the library manager override the parameters that are specified for the library client. If the device class names are different, the library manager uses the parameters that are specified in a device class that matches the device type that is specified for the library client.

7. Define a storage pool for the shared library.
8. Repeat the steps to configure another server as a library client.

-  Example: Library sharing for AIX and Linux servers
To learn how to set up a SCSI library sharing environment for servers that run on AIX® or Linux systems, review the sample procedure.
- Example: Library sharing for Windows servers
To learn how to set up a library sharing environment for servers that run on Windows systems, review the sample procedure.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[DEFINE LIBRARY](#) (Define a library)

[DEFINE STGPOOL](#) (Define a volume in a storage pool)



Example: Library sharing for AIX and Linux servers

To learn how to set up a SCSI library sharing environment for servers that run on AIX® or Linux systems, review the sample procedure.

About this task

In this example, a library manager server named ASTRO and a library client named JUDY are configured. To help clarify where each step is performed, the commands are preceded by the server name from which the command is issued. Most commands are issued from the library client.

For SCSI libraries, define the library by specifying the `libtype=scsi` parameter.

Procedure

1. To set up ASTRO as the library manager server, define a shared SCSI library named SANGROUP. For example:

```
astro> define library sangroup libtype=scsi shared=yes
```

Then complete the rest of the steps as described in Example: Configure a SCSI or virtual tape library with a single drive device type to configure the library.

Tip: You can use the `PERFORM LIBACTION` command to define drives and paths for a library in one step.

2. Define ASTRO as the library manager server by issuing the `DEFINE SERVER` command.

```
judy> define server astro serverpassword=secret hladdress=192.0.2.24  
lladdress=1777 crossdefine=yes
```

3. Define the shared library SANGROUP by issuing the `DEFINE LIBRARY` command. You must use the library manager server name in the `PRIMARYLIBMANAGER` parameter, and use `LIBTYPE=SHARED`.

```
judy> define library sangroup libtype=shared primarylibmanager=astro
```

Ensure that the library name is the same as the library name on the library manager.

4. Define paths from the library manager, ASTRO, to two drives in the shared library by issuing the `DEFINE PATH` command.

AIX

```
astro> define path judy drivea srctype=server desttype=drive  
library=sangroup device=/dev/rmt6  
astro> define path judy driveb srctype=server desttype=drive  
library=sangroup device=/dev/rmt7
```

Linux

```
astro> define path judy drivea srctype=server desttype=drive  
library=sangroup device=/dev/IBMtape6  
astro> define path judy driveb srctype=server desttype=drive  
library=sangroup device=/dev/IBMtape7
```

5. Define all device classes that are associated with the shared library. **AIX**

```
judy> define devclass tape library=sangroup devtype=lto
```

Linux

```
judy> define devclass tape library=sangroup devtype=lto
```

The following parameters for the device class definition must be the same on the library client and the library manager:

- o LIBRARY
- o DRIVEENCRYPTION
- o WORM
- o FORMAT

6. Define a storage pool that is named BACKTAPE for the shared library to use. Issue the `DEFINE STGPOOL` command.

```
judy> define stgpool backtape tape maxscratch=50
```

What to do next

Repeat the procedure to define more library clients to your library manager.

Related reference:

[DEFINE DEVCLASS \(Define a device class\)](#)

Related information:

[DEFINE DRIVE \(Define a drive to a library\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Windows

Example: Library sharing for Windows servers

To learn how to set up a library sharing environment for servers that run on Windows systems, review the sample procedure.

About this task

In this example, a library manager server named ASTRO and a library client named JUDY are configured.

For SCSI libraries, define the library by specifying the `libtype=scsi` parameter.

- **Windows** Setting up the library manager server
You must set up the library manager server in order to configure the IBM Spectrum Protect servers to share SAN-connected devices.
- **Windows** Setting up the library client servers
You must set up one or more library client servers to configure the IBM Spectrum Protect servers to share SAN-connected devices.

Windows

Setting up the library manager server

You must set up the library manager server in order to configure the IBM Spectrum Protect™ servers to share SAN-connected devices.

Procedure

The following procedure is an example of how to set up an IBM Spectrum Protect server that is named ASTRO as a library manager:

1. Ensure that the library manager server is running:
 - a. Start the Windows Services Management Console (`services.msc`).
 - b. Select the service. For example, `TSM Server1`.
 - c. If the service is not running, right-click the service name and click Start.
2. Obtain the library and drive information for the shared library device:
 - a. Run the `tsmdlst.exe` utility. The utility is in the `\Program Files\Tivoli\TSM\server` directory.
3. Define a library whose library type is SCSI. For example:

```
define library sangroup libtype=scsi shared=yes
```

This example uses the default for the library's serial number, which is to have the server obtain the serial number from the library itself at the time that the path is defined. Depending on the capabilities of the library, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device.

4. Define the path from the server to the library.

```
define path astro sangroup srctype=server desttype=library  
device=lb0.0.0.2
```

If you did not include the serial number when you defined the library, the server now queries the library to obtain this information. If you did include the serial number when you defined the library, the server verifies what you defined and issues a message if there is a mismatch.

5. Define the drives in the library.

```
define drive sangroup drivea  
define drive sangroup driveb
```

This example uses the default for the drive's serial number, which is to have the server obtain the serial number from the drive itself at the time that the path is defined. Depending on the capabilities of the drive, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device.

This example also uses the default for the drive's element address, which is to have the server obtain the element number from the drive itself at the time that the path is defined.

The element address is a number that indicates the physical location of a drive within an automated library. The server needs the element address to connect the physical location of the drive to the drive's SCSI address. You can have the server obtain the element number from the drive itself at the time that the path is defined, or you can specify the element number when you define the drive.

Depending on the capabilities of the library, the server might not be able to automatically detect the element address. In this case, you must supply the element address when you define the drive. Element numbers for many libraries are available at IBM® Support Portal for IBM Spectrum Protect.

6. Define the path from the server to each of the drives.

```
define path astro drivea srctype=server desttype=drive library=sangroup
device=mt0.1.0.2
define path astro driveb srctype=server desttype=drive library=sangroup
device=mt0.2.0.2
```

If you did not include the serial number or element address when you defined the drive, the server now queries the drive or the library to obtain this information.

7. Define at least one device class.

```
define devclass tape devtype=dlt library=sangroup
```

8. Check in the library inventory. The following example checks all volumes into the library inventory as scratch volumes. The server uses the name on the bar code label as the volume name.

```
checkin libvolume sangroup search=yes status=scratch
checklabel=barcode
```

9. Set up a storage pool for the shared library with a maximum of 50 scratch volumes.

```
define stgpool backtape tape
description='storage pool for shared sangroup' maxscratch=50
```

Related reference:

- [CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)
- [DEFINE DEVCLASS \(Define a device class\)](#)

Related information:

- [DEFINE DRIVE \(Define a drive to a library\)](#)
- [DEFINE LIBRARY \(Define a library\)](#)
- [DEFINE PATH \(Define a path\)](#)
- [DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Windows

Setting up the library client servers

You must set up one or more library client servers to configure the IBM Spectrum Protect™ servers to share SAN-connected devices.

Before you begin

Ensure that a library manager server is defined.

About this task

You must define the library manager server. Use the following procedure as an example of how to set up an IBM Spectrum Protect server that is named JUDY as a library client.

Procedure

1. Ensure that the library manager server is running:
 - a. Start the Windows Services Management Console (services.msc).

- b. Select the service. For example, TSM Server1.
 - c. If the service is not running, right-click and select Start.
2. Obtain the library and drive information for the shared library device:
 - a. Run the `tsmdlst.exe` utility. The utility is in the `\Program Files\Tivoli\TSM\server` directory.
3. Define the shared library, SANGROUP, and identify the library manager. Ensure that the library name is the same as the library name on the library manager.

```
define library sangroup libtype=shared primarylibmanager=astro
```

4. Define the paths from the library client server to each of the drives by issuing commands on the administrative client:

```
define path judy drivea srctype=server desttype=drive library=sangroup
device=mt0.1.0.3
define path judy driveb srctype=server desttype=drive library=sangroup
device=mt0.2.0.3
```

5. Define at least one device class by issuing commands from the library client:

```
define devclass tape devtype=dlt mountretention=1 mountwait=10
library=sangroup
```

Set the parameters for the device class the same on the library client as on the library manager. Making the device class names the same on both servers is also a good practice, but is not required.

The device class parameters that are specified on the library manager server override those specified for the library client. This is true whether or not the device class names are the same on both servers. If the device class names are different, the library manager uses the parameters specified in a device class that matches the device type specified for the library client.

If a library client requires a setting that is different from what is specified in the library manager's device class (for example, a different mount limit), complete the following steps:

- a. Create an additional device class on the library manager server. Specify the parameter settings that you want the library client to use.
 - b. Create a device class on the library client with the same name and device type as the new device class you created on the library server.
6. Define the storage pool, BACKTAPE, that will use the shared library:

```
define stgpool backtape tape
description='storage pool for shared sangroup' maxscratch=50
```

7. Repeat this procedure to define additional servers as library clients.

Related reference:

[DEFINE DEVCLASS \(Define a device class\)](#)

Related information:

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Setting up a storage pool hierarchy

As part of the implementation process, you must set up a storage pool hierarchy. Set up at least one primary storage pool on disk and one primary storage pool on tape. Ensure that data is migrated from disk to tape daily.

Before you begin

1. Ensure that you reviewed the information in Planning the storage pool hierarchy.
2. Ensure that appropriate rules, also known as *policies*, are specified for backing up client data. Follow the instructions in Specifying rules for backing up and archiving client data.
3. Ensure that a policy is assigned to each node. For instructions about assigning a policy when you register a node, see Registering clients.

Procedure

To set up a storage pool hierarchy, complete the following steps:

1. Define a primary storage pool for the tape device by issuing the DEFINE STGPOOL command.

For example, define a primary storage pool, TAPE1, with a device class of LTO, and enable group collocation. Set the maximum number of scratch volumes that the server can request for this storage pool to 999. Issue the following command:

```
define stgpool tape1 lto pooltype=primary collocate=group
maxscratch=999
```

2. Define the drives, paths, and libraries for the primary storage pool on tape. Follow the instructions in Defining tape devices.
3. Define a primary storage pool for the disk device by issuing the DEFINE STGPOOL command.

For example, define a storage pool, DISK1, with a device class of FILE. Ensure that data can be migrated to the tape storage pool, TAPE1, but prevent automatic migration by specifying 100 for the HIGHMIG parameter and 0 for the LOWMIG parameter. Prevent reclamation by specifying 100 for the RECLAIM parameter. Enable node collocation. Set the maximum number of scratch volumes that the server can request for this storage pool to 9999. Use the MIGPROCESS parameter to specify the number of migration processes. The value of the MIGPROCESS parameter should equal the number of drives in the library minus the number of drives that are reserved for restore operations. Issue the following command:

```
define stgpool disk1 file pooltype=primary nextstgpool=tape1
highmig=100 lowmig=0 reclaim=100 collocate=node maxscratch=9999 migprocess=5
```

For more information about how to set up migration from disk to tape, see Migrating disk storage pools.

What to do next

A storage pool hierarchy includes only primary storage pools. After you set up the storage pool hierarchy, complete the following steps:

1. Create a copy storage pool on a tape device. For instructions, see DEFINE STGPOOL (Define a copy storage pool assigned to sequential access devices).
2. Back up the tape-based primary storage pool to the copy storage pool by using the BACKUP STGPOOL command. For instructions, see BACKUP STGPOOL (Back up primary storage pool data to a copy storage pool).
3. To ensure that data can be recovered in a disaster, set up a procedure for moving tape volumes from the copy storage pool to an offsite location. For instructions, see Preparing for and recovering from a disaster by using DRM.

Related reference:

[CHECKIN LIBVOLUME](#) (Check a storage volume into a library)

Related information:

[DEFINE STGPOOL](#) (Define a volume in a storage pool)

Protecting applications and systems

The server protects data for clients, which can include applications, virtual machines, and systems.

- Adding clients
Following the successful setup of your IBM Spectrum Protect™ server, install and configure client software to begin backing up data.

Configuring LAN-free data movement

You can configure the IBM Spectrum Protect™ client and server so that the client, through a storage agent, can move data directly to storage on a SAN. This function, called LAN-free data movement, is provided by the IBM Spectrum Protect for SAN product.

Procedure

To configure LAN-free data movement, complete the following steps. For details, see the documentation for IBM Spectrum Protect for SAN.

1. Verify the network connection.
2. Establish communications among the client, storage agent, and the server.
3. Install and configure software on client systems.
4. Configure devices on the server for the storage agent to access.

5. Configure IBM Spectrum Protect policies for LAN-free data movement for the client.
6. If you are using shared FILE storage, install and configure IBM® TotalStorage SAN File System or IBM Spectrum Scale™.
 - Windows** Restriction: If an IBM Spectrum Scale volume is formatted by an AIX® server, the Windows system uses TCP/IP to transfer data and not the storage area network.
7. Define paths from the storage agent to drives.
8. Start the storage agent and verify the LAN-free configuration.

What to do next

To help you tune the use of your LAN and SAN resources, you can control the path that data transfers take for clients with the capability of LAN-free data movement. Control the path by using the UPDATE NODE command. For each client, you can select one of the following settings for data read and write operations. Specify data read operations by using the DATAREADPATH parameter and data write operations by using the DATAWRITEPATH parameter. The parameter is optional. The default value is ANY.

LAN (LAN path only)

Specify the LAN value if either of the following conditions is true:

- You want to back up or restore a small amount of data.
- The client does not have SAN connectivity.

LANFREE (LAN-free path only)

Specify the LANFREE value if the client and server are on the same SAN, and any of the following conditions are true:

- You want to back up or restore a large amount of data.
- You want to offload the server processing load to the client.
- You want to relieve LAN congestion.

ANY (Any available path)

A LAN-free path is used if one is available. If a LAN-free path is unavailable, the data is moved by using the LAN.

- Validating your LAN-free configuration
After you configure an IBM Spectrum Protect client for LAN-free data movement, you can verify the configuration and server definitions by using the VALIDATE LANFREE command.

Tape encryption methods

Deciding on the encryption method to use depends on how you want to manage your data.

It is critical to secure client data, especially when that data is sensitive. To ensure that data in onsite and offsite volumes is protected, IBM tape encryption technology is available.

IBM tape technology supports different methods of drive encryption for the following devices:

- IBM 3592 Generation 2 and Generation 3
- IBM Linear Tape-Open (LTO) Generation 4 and Generation 5

The methods of drive encryption that you can use with IBM Spectrum Protect™ are set up at the hardware level. IBM Spectrum Protect cannot control or change which encryption method is used in the hardware configuration. If the hardware is set up for the Application method, IBM Spectrum Protect can turn encryption on or off depending on the DRIVEENCRYPTION value on the device class.

To encrypt all data in a particular logical library or to encrypt data on more than just storage pool volumes, use the Library or System method. If the encryption key manager is set up to share keys, the Library and System methods can share the encryption key, which allows the two methods to be interchanged. IBM Spectrum Protect cannot share or use encryption keys between the Application method and either the Library or the System methods of encryption.

Table 1. Encryption methods

Encryption method	Description
-------------------	-------------

Encryption method	Description
Application encryption	<p>With application-managed encryption, you can create dedicated storage pools that contain encrypted volumes only. This way, you can use storage pool hierarchies and policies to manage the way data is encrypted.</p> <p>Encryption keys are managed by the application, in this case, IBM Spectrum Protect. IBM Spectrum Protect generates and stores the keys in the server database. Data is encrypted during write operations, when the encryption key is passed from the server to the drive. Data is decrypted for read operations.</p> <p>To encrypt storage pool volumes and eliminate some encryption processing on your system, enable the Application method. Use application-managed encryption only for storage pool volumes. Other volumes, such as backup-set tapes, export volumes, and database backups, are not encrypted by using the Application method.</p> <p>Requirement: When application encryption is enabled, you must take extra care to secure database backups because the encryption keys that are used to encrypt and decrypt data are stored in the server database. To restore your data, you must have the correct database backup and corresponding encryption keys to access your information. Ensure that you back up the database frequently and safeguard the backups to prevent data loss or theft. Anyone who has access to both the database backup and the encryption keys has access to your data.</p>
Library encryption	<p>With library-managed encryption, you can control which volumes are encrypted by using their serial numbers. You can specify a range or set of volumes to encrypt.</p> <p>Encryption keys are managed by the library. Keys are stored in an encryption key manager and provided to the drive. If you set up the hardware to use library-managed encryption, you can use this method by issuing the DEFINE DEVCLASS command and specifying the DRIVEENCRYPTION=ALLOW parameter.</p> <p>Restriction: Only certain IBM libraries support IBM LTO-4 and later encryption. For details, see Configuring tape drive encryption.</p>
System encryption	<p>System-managed encryption is available only on the AIX® operating system. Encryption keys that are provided to the drive are managed by the device driver or operating system and stored in an encryption key manager. If the hardware is set up to use system encryption, you can use this method by issuing the DEFINE DEVCLASS command and specifying the DRIVEENCRYPTION=ALLOW parameter.</p>

To determine whether a volume is encrypted and which method was used, issue the QUERY VOLUME command and specify the FORMAT=DETAILED parameter.

- **Configuring tape drive encryption**
You can use drive encryption to protect tapes that contain critical or sensitive data, for example, tapes that contain confidential financial information. Drive encryption can be useful when you move tapes from the IBM Spectrum Protect server environment to an onsite or offsite location.

Controlling tape storage operations

Device class definitions for tapes include parameters that allow you to control storage operations.

- **How IBM Spectrum Protect fills volumes**
The DEFINE DEVCLASS command has an optional ESTCAPACITY parameter that indicates the estimated capacity for sequential volumes that are associated with the device class. IBM Spectrum Protect™ uses the estimated capacity of volumes to determine the estimated capacity of a storage pool, and the estimated percent utilized.
- **Specifying the estimated capacity of tape volumes**
IBM Spectrum Protect also uses estimated capacity to determine when to begin the reclamation of storage pool volumes.
- **Specifying recording formats for tape media**
You can specify the recording format that is used by IBM Spectrum Protect to write data to tape media. If you plan to mix generations of drives, or different drive types, within a library, you must specify a recording format for each drive generation and each drive type. In this way, the server can differentiate between the drive generations and drive types.
- **Associating library objects with device classes**
A library contains the drives that can be used to mount the volume. Only one library can be associated with a device class. However, multiple device classes can reference the same library.
- **Controlling media-mount operations for tape devices**
By using device class definitions, you can control the number of mounted volumes, the amount of time a volume remains mounted, and the amount of time that the IBM Spectrum Protect server waits for a drive to become available.
- **Preempting operations**
The server can preempt server or client operations for a higher priority operation when a mount point is in use and no others are available, or access to a specific volume is required. When an operation is preempted, it is canceled.
- **Impacts of device changes on the SAN**
The SAN environment can shift dramatically due to device or cabling changes. The dynamic nature of the SAN can cause static definitions to fail or become unpredictable.
- **Windows** **Displaying device information**
You can display information about devices that are connected to the server by using the device information utility (tsmdlst).
- **Write-once, read-many tape media**
Write-once, read-many (WORM) media help to prevent accidental or deliberate deletion of critical data. However, IBM Spectrum Protect imposes certain restrictions and guidelines to follow when you use WORM media.
- **Windows** **Troubleshooting problems with devices**
You can troubleshoot errors that occur when you configure or use devices with IBM Spectrum Protect.

How IBM Spectrum Protect fills volumes

The DEFINE DEVCLASS command has an optional ESTCAPACITY parameter that indicates the estimated capacity for sequential volumes that are associated with the device class. IBM Spectrum Protect™ uses the estimated capacity of volumes to determine the estimated capacity of a storage pool, and the estimated percent utilized.

If the ESTCAPACITY parameter is not specified, IBM Spectrum Protect uses a default value that is based on the recording format that is specified for the device class by using the FORMAT parameter.

If you specify an estimated capacity that exceeds the actual capacity of the volume in the device class, IBM Spectrum Protect updates the estimated capacity of the volume when the volume becomes full. When IBM Spectrum Protect reaches the end of the volume, it updates the capacity to match the amount that is written to the volume.

You can either accept the default estimated capacity for the device class, or explicitly specify an estimated capacity. An accurate estimated capacity value is not required, but is useful. IBM Spectrum Protect uses the estimated capacity of volumes to determine the estimated capacity of a storage pool, and the estimated percent that is used. You might want to change the estimated capacity if on or both of the following conditions are true:

- The default estimated capacity is inaccurate because of data compression.
- You have volumes of nonstandard size.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Specifying the estimated capacity of tape volumes

IBM Spectrum Protect™ also uses estimated capacity to determine when to begin the reclamation of storage pool volumes.

About this task

For tape device classes, the default values selected by the server depend on the recording format that is used to write data to the volume. You can either accept the default for a device type or specify a value.

To specify estimated capacity for tape volumes, use the ESTCAPACITY parameter when you define the device class or update its definition.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Specifying recording formats for tape media

You can specify the recording format that is used by IBM Spectrum Protect™ to write data to tape media. If you plan to mix generations of drives, or different drive types, within a library, you must specify a recording format for each drive generation and each drive type. In this way, the server can differentiate between the drive generations and drive types.

About this task

To specify a recording format, use the FORMAT parameter when you define the device class or update its definition.

If all drives associated with that device class are identical, specify FORMAT=DRIVE. The server selects the highest format that is supported by the drive on which a volume is mounted.

If some drives associated with the device class support a higher density format than others, specify a format that is compatible with all drives.

If drives in a single SCSI library use different tape technologies (for example, DLT and LTO Ultrium), specify a unique value for the FORMAT parameter in each device class definition.

For a configuration example, see Example: Configure a SCSI or virtual tape library with multiple drive device types.

The recording format that the server uses for a volume is selected when data is first written to the volume. Updating the FORMAT parameter does not affect media that already contain data until those media are rewritten from the beginning. This process might happen after a volume is reclaimed or deleted, or after all of the data on the volume expires.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Associating library objects with device classes

A library contains the drives that can be used to mount the volume. Only one library can be associated with a device class. However, multiple device classes can reference the same library.

About this task

To associate a device class with a library, use the LIBRARY parameter when you define a device class or update its definition.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Controlling media-mount operations for tape devices

By using device class definitions, you can control the number of mounted volumes, the amount of time a volume remains mounted, and the amount of time that the IBM Spectrum Protect™ server waits for a drive to become available.

- Controlling the number of simultaneously mounted volumes
When you set a mount limit for a device class, you must consider the number of storage devices that are connected to your system. You must also consider whether you use the simultaneous-write function, whether you associate multiple device classes with a single library, and the number of processes that run at the same time.
- Controlling the amount of time that a volume remains mounted
You can control the amount of time that a mounted volume remains mounted after its last I/O activity. If a volume is used frequently, you can improve performance by setting a longer mount retention period to avoid unnecessary mount and dismount operations.
- Controlling the amount of time that the server waits for a drive
You can specify the maximum amount of time, in minutes, that the IBM Spectrum Protect server waits for a drive to become available for the current mount request.

Controlling the number of simultaneously mounted volumes

When you set a mount limit for a device class, you must consider the number of storage devices that are connected to your system. You must also consider whether you use the simultaneous-write function, whether you associate multiple device classes with a single library, and the number of processes that run at the same time.

About this task

When you select a mount limit for a device class, consider the following issues:

- How many storage devices are connected to your system?

Do not specify a mount limit value that is greater than the number of associated available drives in your installation. If the server tries to mount as many volumes as specified by the mount limit and no drives are available for the required volume, an error occurs and client sessions might end. (This restriction does not apply when the DRIVES parameter is specified.)

If you are sharing library resources on a SAN among IBM Spectrum Protect™ servers, you must limit the number of tape drives that a library client can use at a time. To allow multiple library client servers use a library simultaneously specify the MOUNTLIMIT parameter when you define or update the device class on the library client. For more information about configuring library sharing, see [Configuring library sharing](#).

- Are you using the simultaneous-write function to primary storage pools, copy storage pools, and active-data pools?

Specify a mount limit value that provides enough mount points to support writing data simultaneously to the primary storage pool and all associated copy storage pools and active-data pools.

- Are you associating multiple device classes with a single library?

A device class that is associated with a library can use any drive in the library that is compatible with the device class' device type. Because you can associate more than one device class with a library, a single drive in the library can be used by more than one device class. IBM Spectrum Protect ensures that two operations cannot use the same drive simultaneously by using two different device classes.

- How many IBM Spectrum Protect processes do you want to run at the same time, by using devices in this device class?

IBM Spectrum Protect automatically cancels some processes to run other, higher priority processes. If the server is using all available drives in a device class to complete higher priority processes, lower-priority processes must wait until a drive becomes available. For example, IBM Spectrum Protect cancels the process for a client that backs up directly to tape if the drive is needed for a server migration or tape reclamation process. IBM Spectrum Protect cancels a tape reclamation process if the drive is needed for a client restore operation. For more information, see [Preempting operations](#).

If processes are often canceled by other processes, consider whether you can make more drives available for IBM Spectrum Protect use. Otherwise, review your scheduling of operations to reduce the contention for drives.

This consideration also applies to the simultaneous-write function. You must have enough drives available to allow for a successful simultaneous-write operation.

To specify the maximum number of volumes that can be simultaneously mounted, use the MOUNTLIMIT parameter when you define the device class or update its definition.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Controlling the amount of time that a volume remains mounted

You can control the amount of time that a mounted volume remains mounted after its last I/O activity. If a volume is used frequently, you can improve performance by setting a longer mount retention period to avoid unnecessary mount and dismount operations.

About this task

If mount operations are being handled by manual, operator-assisted activities, you might want to specify a long mount retention period. For example, if only one operator supports your entire operation on a weekend, then define a long mount retention period so that the operator is not being asked to mount volumes every few minutes.

To control the amount of time a mounted volume remains mounted, use the MOUNTRETENTION parameter when you define the device class or update its definition. For example, if the mount retention value is 60, and a mounted volume remains idle for 60 minutes, the server dismounts the volume.

While IBM Spectrum Protect™ has a volume mounted, the drive is allocated to IBM Spectrum Protect and cannot be used for anything else. If you need to free the drive for other uses, you can cancel IBM Spectrum Protect operations that are using the drive and then dismount the volume. For example, you can cancel server migration or backup operations. For information on how to cancel processes and dismount volumes, see [Managing server requests for volumes](#)

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Controlling the amount of time that the server waits for a drive

You can specify the maximum amount of time, in minutes, that the IBM Spectrum Protect™ server waits for a drive to become available for the current mount request.

About this task

To control the wait time for a drive to become available for a mount request, use the MOUNTWAIT parameter when you define or update a device class.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Preempting operations

The server can preempt server or client operations for a higher priority operation when a mount point is in use and no others are available, or access to a specific volume is required. When an operation is preempted, it is canceled.

You can use the QUERY MOUNT command to see the status of the volume for the mount point.

By default, preemption is enabled on the server. To disable preemption, specify the NOPREEMPT option in the server options file. If you specify this option, the BACKUP DB command, and the export and import commands are the only operations that can preempt other operations.

- Mount point preemption
If a high-priority operation requires a mount point that is in a specific device class and all the mount points in the device class are in use, the high-priority operation can preempt a mount point from a lower-priority operation.
- Volume access preemption
If a high-priority operation requires access to a specific volume and that volume is in use, the high-priority operation can preempt the lower-priority operation for that volume.

Related reference:

[BACKUP DB \(Back up the database\)](#)

[QUERY MOUNT \(Display information on mounted sequential access volumes\)](#)

Mount point preemption

If a high-priority operation requires a mount point that is in a specific device class and all the mount points in the device class are in use, the high-priority operation can preempt a mount point from a lower-priority operation.

Mount points can be preempted only when the device class of the operation preempting and the operation that is being preempted is the same.

The following high-priority operations can preempt other operations for a mount point.

- Database backup operations
- Retrieve, restore, or HSM recall operations that are initiated by clients
- Restore operations by using a remote data mover
- Export operations
- Import operations
- Operations to generate backup sets

The following server operations cannot preempt other operations or be preempted:

- Audit a volume
- Restore data from a copy or active-data pool
- Prepare a recovery plan file
- Store data by using a remote data mover

The following operations can be preempted and are listed in order of priority, from highest priority to lowest priority. The server selects the lowest priority operation to preempt, for example, identify duplicates.

- Replicate nodes
- Back up data to a copy storage pool
- Copy active data to an active data pool
- Move data on a storage pool volume
- Migrate data from disk to sequential media
- Migrate data from sequential media to sequential media
- Back up, archive, or HSM migrate operations that are initiated by clients
- Reclaim volumes in a sequential-access storage pool
- Identify duplicates

Volume access preemption

If a high-priority operation requires access to a specific volume and that volume is in use, the high-priority operation can preempt the lower-priority operation for that volume.

For example, if a restore request requires access to a volume in use by a reclamation operation and a drive is available, the reclamation operation is canceled.

The following high-priority operations can preempt operations for access to a specific volume:

- Database backup operations
- Retrieve, restore, or HSM recall operations that are initiated by clients
- Restore operations by using a remote data mover
- Export operations
- Import operations
- Operations to generate backup sets

The following operations cannot preempt other operations or be preempted:

- Audit volume
- Restore data from a copy or active-data pool
- Prepare a recovery plan

- Store data by using a remote data mover

The following operations can be preempted, and are listed in order of priority, from highest priority to lowest priority. The server selects the lowest priority operation to preempt, for example, identify duplicates.

- Replicate nodes
- Back up data to a copy storage pool
- Copy active data to an active data pool
- Move data on a storage pool volume
- Migrate data from disk to sequential media
- Migrate data from sequential media to sequential media
- Back up, archive, or HSM migrate data that is initiated by client
- Reclaim volumes in a sequential-access storage pool
- Identify duplicates

Impacts of device changes on the SAN

The SAN environment can shift dramatically due to device or cabling changes. The dynamic nature of the SAN can cause static definitions to fail or become unpredictable.

Device IDs that are assigned by the SAN and known to the server or storage agent can be altered due to bus resets or other environmental changes. For example, the server might know a device X as *rmt0* (on AIX®), based on the original path specification to the server and original configuration of the LAN. However, some event in the SAN, for example, the addition of new device Y, causes device X to be assigned *rmt1*. When the server tries to access device X by using *rmt0*, either the access fails or the wrong target device is accessed. The server attempts to recover from changes to devices on the SAN by using device serial numbers to confirm the identity of devices it contacts.

When you define a drive or library, you have the option of specifying the serial number for that device. If you do not specify the serial number when you define the device, the server obtains the serial number when you define the path for the device. In either case, the server then has the device serial number in its database and can use it to confirm the identity of a device for operations.

When the server uses drives and libraries on a SAN, the server attempts to verify that the correct device is used. The server contacts the device by using the device name in the path that you defined for it. The server then requests the serial number from the device, and compares that serial number with the serial number that is stored in the server database for that device.

If the serial number does not match, the server begins the process of discovery on the SAN, attempting to find the device with the matching serial number. If the server finds the device with the matching serial number, it corrects the definition of the path in the server's database by updating the device name in that path. The server issues a message with information about the change that is made to the device. Then, the server proceeds to use the device.

To determine when device changes on the SAN affect the IBM Spectrum Protect™ server, you can monitor the activity log for messages. The following messages are related to serial numbers:

- ANR8952 through ANR8958
- ANR8961 through ANR8968
- ANR8974 through ANR8975

Restriction: Some devices cannot report their serial numbers to applications such as the IBM Spectrum Protect server. If the server cannot obtain the serial number from a device, the server cannot help the system to recover from a device location change on the SAN.

Windows

Displaying device information

You can display information about devices that are connected to the server by using the device information utility (tsmdlst).

Before you begin

- Ensure that the HBA API is installed. The HBA API is required to run the device information utility.
- Ensure that the tape device driver is installed and configured.

Procedure

1. From a command prompt, change to the `server` subdirectory of the server installation directory, for example, `C:\Program Files\Tivoli\TSM\server`.
2. Run the `tsmdlst.exe` executable file.

Related reference:

- [QUERY SAN](#) (Query the devices on the SAN)
- [tsmdlst](#) (Display information about devices)

Write-once, read-many tape media

Write-once, read-many (WORM) media help to prevent accidental or deliberate deletion of critical data. However, IBM Spectrum Protect™ imposes certain restrictions and guidelines to follow when you use WORM media.

You can use the following types of WORM media with IBM Spectrum Protect:

- IBM® 3592, all supported generations
- IBM LTO-3 and all supported generations
- HP LTO-3 and all supported generations
- Quantum LTO-3 and all supported generations
- Quantum SDLT 600, Quantum DLT V4, and Quantum DLT S4
- StorageTek VolSafe
- Sony AIT50 and AIT100

Tips:

- A storage pool can consist of either WORM or RW media, but not both.
- To avoid wasting tape after a restore or import operation, do not use WORM tapes for database backup or export operations.
- WORM-capable drives
To use WORM media in a library, all the drives in the library must be WORM-capable. A mount will fail if a WORM cartridge is mounted in a read/write (RW) drive.
- Check-in of WORM media
The type of WORM media determines whether the media label needs to be read during check-in.
- Restrictions on WORM media
You cannot use prelabeled WORM media with the LTO or ECARTRIDGE device class.
- Mount failures with WORM media
If WORM tape media are loaded into a drive for a read-write (RW) device-class mount, it will cause a mount failure. Similarly, if RW tape media are loaded into a drive for a WORM device-class mount, the mount will fail.
- Relabeling WORM media
You cannot relabel a WORM cartridge if it contains data. This applies to Sony AIT WORM, LTO WORM, SDLT WORM, DLT WORM, and IBM 3592 cartridges. The label on a VolSafe volume should be overwritten only once and only if the volume does not contain usable, deleted, or expired data.
- Removing private WORM volumes from a library
If you perform an action on a WORM volume (for example, if you delete file spaces) and the server does not mark the volume as full, the volume is returned to scratch status. If a WORM volume is not marked as full and you delete it from a storage pool, the volume remains private. To remove a private WORM volume from a library, you must issue the CHECKOUT LIBVOLUME command.
- Creation of DLT WORM volumes
DLT WORM volumes can be converted from read/write (RW) volumes.
- Support for short and normal 3592 WORM tapes
IBM Spectrum Protect supports both short and normal 3592 WORM tapes. For best results, define them in separate storage pools
- Querying a device class for the WORM-parameter setting
You can determine the setting of the WORM parameter for a device class by using the QUERY DEVCLASS command. The output contains a field, labeled WORM, and a value (YES or NO).

WORM-capable drives

To use WORM media in a library, all the drives in the library must be WORM-capable. A mount will fail if a WORM cartridge is mounted in a read/write (RW) drive.

However, a WORM-capable drive can be used as a RW drive if the WORM parameter in the device class is set to NO. Any type of library can have both WORM and RW media if *all* of the drives are WORM enabled. The only exception to this rule is NAS-attached libraries in which WORM tape media cannot be used.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

Related information:

[UPDATE DEVCLASS](#) (Update a device class)

Check-in of WORM media

The type of WORM media determines whether the media label needs to be read during check-in.

Library changers cannot identify the difference between standard read/write (RW) tape media and the following types of WORM tape media:

- VolSafe
- Sony AIT
- LTO
- SDLT
- DLT

To determine the type of WORM media that is being used, a volume must be loaded into a drive. Therefore, when you check in one of these types of WORM volumes, you must use the CHECKLABEL=YES option on the CHECKIN LIBVOLUME command.

If they provide support for WORM media, IBM® 3592 library changers can detect whether a volume is WORM media without loading the volume into a drive. Specifying CHECKLABEL=YES is not required. Verify with your hardware vendors that your 3592 drives and libraries provide the required support.

Related reference:

[CHECKIN LIBVOLUME](#) (Check a storage volume into a library)

Restrictions on WORM media

You cannot use prelabeled WORM media with the LTO or ECARTRIDGE device class.

You cannot use WORM media with IBM Spectrum Protect™ specified as the drive-encryption key manager for the following drives:

- IBM® LTO-5, LTO-6, and later
- HP LTO-5, LTO-6, and later
- Oracle StorageTek T10000B
- Oracle StorageTek T10000C
- Oracle StorageTek T10000D

Mount failures with WORM media

If WORM tape media are loaded into a drive for a read-write (RW) device-class mount, it will cause a mount failure. Similarly, if RW tape media are loaded into a drive for a WORM device-class mount, the mount will fail.

Relabeling WORM media

You cannot relabel a WORM cartridge if it contains data. This applies to Sony AIT WORM, LTO WORM, SDLT WORM, DLT WORM, and IBM® 3592 cartridges. The label on a VolSafe volume should be overwritten only once and only if the volume does not contain usable, deleted, or expired data.

Issue the LABEL LIBVOLUME command only once for VolSafe volumes. You can guard against overwriting the label by using the OVERWRITE=NO option on the LABEL LIBVOLUME command.

Related reference:

[LABEL LIBVOLUME](#) (Label a library volume)

Removing private WORM volumes from a library

If you perform an action on a WORM volume (for example, if you delete file spaces) and the server does not mark the volume as full, the volume is returned to scratch status. If a WORM volume is not marked as full and you delete it from a storage pool, the volume remains private. To remove a private WORM volume from a library, you must issue the CHECKOUT LIBVOLUME command.

Related reference:

[CHECKOUT LIBVOLUME](#) (Check a storage volume out of a library)

Creation of DLT WORM volumes

DLT WORM volumes can be converted from read/write (RW) volumes.

If you have SDLT-600, DLT-V4, or DLT-S4 drives and you want to enable them for WORM media, upgrade the drives by using V30 or later firmware available from Quantum. You can also use DLTIce software to convert unformatted RW volumes or blank volumes to WORM volumes.

In SCSI libraries, the IBM Spectrum Protect™ server creates scratch DLT WORM volumes automatically when the server cannot locate any scratch WORM volumes in a library's inventory. The server converts available unformatted or blank RW scratch volumes or empty RW private volumes to scratch WORM volumes. The server also rewrites labels on newly created WORM volumes by using the label information on the existing RW volumes.

Support for short and normal 3592 WORM tapes

IBM Spectrum Protect™ supports both short and normal 3592 WORM tapes. For best results, define them in separate storage pools

Querying a device class for the WORM-parameter setting

You can determine the setting of the WORM parameter for a device class by using the QUERY DEVCLASS command. The output contains a field, labeled WORM, and a value (YES or NO).

Related information:

[QUERY DEVCLASS](#) (Display information on one or more device classes)

Windows

Troubleshooting problems with devices

You can troubleshoot errors that occur when you configure or use devices with IBM Spectrum Protect™.

About this task

Use Table 1 to find a solution to the device-related problem.

Table 1. Resolving device problems

Symptom	Problem	Solution			
Conflicts with other applications.	IBM Spectrum Protect requires a storage area network to share devices.	Set up a storage area network. Attention: Data loss can occur if multiple IBM Spectrum Protect servers use the same device. Define or use a device with only one IBM Spectrum Protect server. <table border="1"><tr><td>AIX</td><td>Linux</td><td>Other applications can access IBM Spectrum Protect devices, by using a SCSI tape driver.</td></tr></table>	AIX	Linux	Other applications can access IBM Spectrum Protect devices, by using a SCSI tape driver.
AIX	Linux	Other applications can access IBM Spectrum Protect devices, by using a SCSI tape driver.			

Symptom	Problem	Solution
Labeling fails.	A device for labeling volumes cannot be used at the same time that the server uses the device for other processes.	You cannot overwrite existing volumes in a storage pool. You must resolve any hardware issues before you label a volume.
	Incorrect or incomplete license registration.	Register the license for the device support that was purchased.
Conflicts among device drivers.	IBM Spectrum Protect issues messages about I/O errors when you define or use a sequential access device.	Windows device drivers and drivers that are provided by other applications can interfere with the IBM Spectrum Protect device driver if the IBM Spectrum Protect driver is not started first. To check on the order that device drivers are started by the system, complete the following steps: <ol style="list-style-type: none"> 1. Click Control Panel. 2. Click Devices. Device drivers and their startup types are listed.
I/O errors	When you try to define or use a tape device, there might be device-driver conflicts. Windows device drivers and drivers that are provided by other applications can interfere with the IBM Spectrum Protect device driver if it is not started first.	

Completing the implementation

After the IBM Spectrum Protect™ solution is configured and running, test backup operations and set up monitoring to ensure that everything runs smoothly.

Procedure

1. Test backup operations to verify that your data is protected in the way that you expect.
 - a. On the Clients page of the Operations Center, select the clients that you want to back up, and click Back Up.
 - b. On the Servers page of the Operations Center, select the server for which you want to back up the database. Click Back Up, and follow the instructions in the Back Up Database window.
 - c. Verify that the backup operations completed successfully with no warning or error messages.
Tip: Alternatively, you can use the backup-archive client GUI to back up client data and you can backup the server database by issuing BACKUP DB command from an administrative command-line.
2. Set up monitoring for your solution by following the instructions in Monitoring a tape solution.

Monitoring a tape solution

After you implement an IBM Spectrum Protect™ tape-based solution, monitor the solution to ensure correct operation. By monitoring the solution daily and periodically, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

About this task

The preferred way to monitor a solution is by using the Operations Center, which provides overall and detailed system status in a graphical user interface. In addition, you can configure the Operations Center to generate email reports that summarize system status.

Procedure


1. Complete daily monitoring tasks. For instructions, see Daily monitoring checklist.
2. Complete periodic monitoring tasks. For instructions, see Periodic monitoring checklist.
3. Verify that your system complies with licensing requirements. For instructions, see Verifying license compliance.
4. Optional: Set up email reports of system status. For instructions, see Tracking system status by using email reports


Daily monitoring checklist

To ensure that you are completing the daily monitoring tasks for your IBM Spectrum Protect™ solution, review the daily monitoring checklist.

Complete the daily monitoring tasks from the Operations Center Overview page. You can access the Overview page by opening the Operations Center and clicking Overviews.

The following figure shows the location for completing each task.







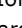
Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.


The following table lists the daily monitoring tasks and provides instructions for completing each task.

Table 1. Daily monitoring tasks



Task	Basic procedures	Advanced procedures and troubleshooting information
------	------------------	---



Task	Basic procedures	Advanced procedures and troubleshooting information
<p>Watch for security notifications, which can indicate a ransomware attack.</p>	<p>If a potential ransomware attack is detected in the IBM Spectrum Protect environment, a security notification message is displayed in the foreground of the Operations Center. For more information, click the message to open the Security Notifications page.</p>	<p>On the Security Notifications page, you can take the following actions:</p> <ul style="list-style-type: none"> • View notification details by client. Restriction: In Operations Center Version 8.1.5, notifications are available only for backup-archive clients. • Acknowledge a security notification by selecting it and clicking Acknowledge. When you acknowledge a security notification, a check mark is added to the Acknowledged column of the Security Notifications page for the selected client. The standard by which a notification is acknowledged is determined by your organization. A check mark might mean that you investigated the issue and determined that it is a false positive. Or it might mean that a problem exists and is being resolved. • Assign a security notification to an administrator by selecting the security notification and clicking Assign. To view the assignment, the administrator must sign in to the Operations Center and click Overviews > Security. If you are not certain whether the administrator regularly monitors the Security Notifications page, notify the administrator about the assignment. • If the notification is a false positive, you can select the security notification and click Reset. The security notification is deleted. Historical data that is used for baseline comparisons with the most recent backup operation is deleted. A new baseline is calculated going forward.
<p>1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.</p>	<p>To verify whether clients are at risk, in the Clients area, look for an At risk notification. To view details, click the Clients area. Attention: If the At risk percentage is much greater than usual, it might indicate a ransomware attack. A ransomware attack can cause backup operations to fail, thus placing clients at risk. For example, if the percentage of clients at risk is normally between 5% and 10%, but the percentage increases to 40% or 50%, investigate the cause. If you installed the client management service on a backup-archive client, you can view and analyze the client error and schedule logs by completing the following steps:</p> <ol style="list-style-type: none"> 1. In the Clients table, select the client and click Details. 2. To diagnose an issue, click Diagnosis. 	<p>For clients that do not have the client management service installed, access the client system to review the client error logs.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>2 Determine whether client-related or server-related errors require attention.</p>	<p>To determine the severity of any reported alerts, in the Alerts area, hover over the columns.</p>	<p>To view additional information about alerts, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Alerts area. 2. In the Alerts table, select an alert. 3. In the Activity Log pane, review the messages. The pane displays related messages that were issued before and after the selected alert occurred.
<p>3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.</p>	<ol style="list-style-type: none"> 1. To verify whether servers are at risk, in the Servers area, look for an Unavailable notification. 2. To view additional information, click the Servers area. 3. Select a server in the Servers table and click Details. 	<p>Tip: If you detect an issue that is related to server properties, update the server properties:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a server and click Details. 2. To update server properties, click Properties.
<p>4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.</p>	<ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Status column of the table, view the status of the server and resolve any issues: <ul style="list-style-type: none"> o Normal  Sufficient space is available for the server database, active log, and archive log. o Critical  Insufficient space is available for the server database, active log, or archive log. You must add space immediately, or the data protection services that are provided by the server will be interrupted. o Warning  The server database, active log, or archive log is running out of space. If this condition persists, you must add space. o Unavailable  Status cannot be obtained. Ensure that the server is running, and that there are no network issues. This status is also shown if the monitoring administrator ID is locked or otherwise unavailable on the server. This ID is named IBM-OC-hub_server_name. o Unmonitored  Unmonitored servers are defined to the hub server, but are not configured for management by the Operations Center. To configure an unmonitored server, select the server, and click Monitor Spoke. 	<p>You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>5 Verify server database backup operations.</p>	<p>To determine when a server was most recently backed up, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column. 	<p>To obtain more detailed information about backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a row and click Details. 2. In the DB Backup area, hover over the check marks to review information about backup operations. <p>If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click Back Up. <p>To determine whether the server database is configured for automatic backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the menu bar, hover over the settings icon  and click Command Builder. 2. Issue the QUERY DB command: <pre>query db f=d</pre> 3. In the output, review the Full Device Class Name field. If a device class is specified, the server is configured for automatic database backups.
<p>6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.</p>	<p>To search for information about processes that failed because of server issues, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click Servers > Maintenance. 2. To obtain the two-week history of a process, view the History column. 3. To obtain more information about a scheduled process, hover over the check box that is associated with the process. 	<p>For more information about monitoring processes and resolving issues, see the Operations Center online help.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>7 Verify that the amount of data that was recently sent to and from servers is within the expected range.</p>	<ul style="list-style-type: none"> To obtain an overview of activity in the last 24 hours, view the Activity area. To compare activity in the last 24 hours with activity in the previous 24 hours, review the figures in the Current and Previous areas. 	<ul style="list-style-type: none"> If more data was sent to the server than you expected, determine which clients are backing up more data and investigate the cause. It is possible that client-side data deduplication is not working correctly. Attention: If the amount of backed-up data is significantly larger than usual, it might indicate a ransomware attack. When ransomware encrypts data, the system perceives the data as being changed, and the changed data is backed up. Thus, backup volumes become larger. To determine which clients are affected, click the Applications, Virtual, or Systems tabs. If less data was sent to the server than you expected, investigate whether client backup operations are proceeding on schedule.

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>8 Verify that storage pools are available to back up client data.</p>	<ol style="list-style-type: none"> 1. If problems are indicated in the Storage & Data Availability area, click Pools to view the details: <ul style="list-style-type: none"> o If the Critical  status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. Attention: If the status is critical, investigate the cause: <ul style="list-style-type: none"> ■ If the data deduplication rate for a storage pool drops significantly, it might indicate a ransomware attack. During a ransomware attack, data is encrypted and cannot be deduplicated. To verify the data deduplication rate, in the Storage Pools table, review the value in the % Savings column. ■ If a storage pool unexpectedly becomes 100% utilized, it might indicate a ransomware attack. To verify the utilization, review the value in the Capacity Used column. Hover over the values to see the percentages of used and free space. o If the Warning  status is displayed, the storage pool is running out of space, or its access status is read-only. 2. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column. 	<p>To view the storage-pool capacity that was used over the past two weeks, select a row in the Storage Pools table and click Details.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>9 Verify that storage devices are available for backup operations.</p>	<p>In the Storage & Data Availability area, in the Volumes section, under the capacity bars, review the status that is reported next to Devices. If a Critical  or Warning  status is displayed for any device, investigate the issue. To view details, click Devices.</p>	<p>Tape devices might have a warning or critical status if drives are unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. A tape device might also have a critical status if the library is offline. Other columns of the Tape Devices table show the state of the library robotics, drives, and paths.</p> <p>To resolve issues with tape drives that have a critical state, you can take the drive offline if you need to use it for another activity, such as maintenance. To take a drive offline, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations CenterStorage page, and select Tape Devices. 2. To view more information about a tape library, select a row and click Details. 3. To take a drive offline, select the tape drive and click Offline. <p>For tape backup operations, verify that sufficient scratch tapes are available. If you are not certain whether the number of available scratch tapes is sufficient, open the details notebook to view tape usage and an estimate of scratch tape availability. To open the details notebook, select a library in the table and click Details.</p>

Periodic monitoring checklist

To help ensure that operations run correctly, complete the tasks in the periodic monitoring checklist. Schedule periodic tasks frequently enough so that you can detect potential issues before they become problematic.

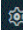
Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.

Table 1. Periodic monitoring tasks

Task	Basic procedures	Advanced procedures and troubleshooting
------	------------------	---

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor system performance.	<p>Determine the length of time that is required for client backup operations:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. Find the server that is associated with the client. 2. Click Servers. Select the server and click Details. 3. To view the duration of completed tasks in the last 24 hours, click Completed Tasks. 4. To view the duration of tasks that were completed more than 24 hours ago, use the QUERY ACTLOG command. For information about this command, see . 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. <p>If you installed the client management service on a backup-archive client, you can diagnose performance issues for the backup-archive client by completing the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. 2. Select a backup-archive client and click Details. 3. To retrieve client logs, click Diagnosis. 	<p>Limit the time for client backup operations to 8 - 12 hours. Ensure that client schedules do not overlap with server maintenance tasks.</p> <p>For instructions about reducing the time that it takes for the client to back up data to the server, see Resolving common client performance problems.</p> <p>Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks.</p> <p>For information about identifying and resolving other performance issues, see Performance.</p>

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Verify that current backup files for device configuration and volume history information are saved.</p>	<p>Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations.</p> <p>To locate the volume history and device configuration files, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon and click Command Builder. 2. To locate the volume history and device configuration files, issue the following commands: <ul style="list-style-type: none"> <code>query option volhistory</code> <code>query option devconfig</code> 3. In the output, review the Option Setting column to find the file locations. <p>If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.</p>	

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Determine whether sufficient space is available in the directory for the server instance.</p>	<p>Verify that at least 50 GB of free space is available in the directory for the server instance. Take the action that is appropriate for your operating system:</p> <ul style="list-style-type: none"> <p>AIX To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -g instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Linux To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -h instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Windows In the Windows Explorer program, right-click the file system and click Properties. View the capacity information.</p> <p>The preferred location of the instance directory depends on the operating system where the server is installed:</p> <ul style="list-style-type: none"> <p>AIX Linux /home/tsminst1/tsminst1</p> <p>Windows C:\tsminst1</p> <p>Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.</p>	
<p>Identify unexpected client activity.</p>	<p>To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps:</p> <ol style="list-style-type: none"> On the Operations Center Overview page, click the Clients area. To view activity over the past two weeks, double-click any client. To view the number of bytes sent to the client, click the Properties tab. In the Last Session area, view the Sent to client row. 	<p>When you double-click a client in the Clients table, the Activity over 2 Weeks area displays the amount of data that the client sent to the server each day.</p> <p>Periodically review the SQL activity summary table, which contains statistics about client sessions. To compare current activity with previous activity, use an SQL SELECT statement. If the level of activity is significantly different from previous activity, it might indicate a ransomware attack.</p> <p>Periodically review the activity log. Look for ANE messages that indicate how many files were backed up and inspected. Compare current data deduplication rates with previous rates. If an unusually high number of files were backed up, or the rate of data deduplication unexpectedly drops to 0, it might indicate a ransomware attack.</p>

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Monitor storage pool growth over time.</p>	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Pools area. 2. To view the capacity that was used over the last two weeks, select a pool and click Details. 	<p>Tips:</p> <ul style="list-style-type: none"> • To specify the time period that must elapse before all deduplicated extents are removed from a directory-container storage pool or cloud-container storage pool after they are no longer referenced by the inventory, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. 3. Specify the duration in the <code>Delay period for container reuse</code> field. • To determine data deduplication performance for directory-container and cloud-container storage pools, use the <code>GENERATE DEDUPSTATS</code> command. • To view data deduplication statistics for a storage pool, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. <p>Alternatively, use the <code>QUERY EXTENTUPDATES</code> command to display information about updates to data extents in directory-container or cloud-container storage pools. The command output can help you determine which data extents are no longer referenced and which ones are eligible to be deleted from the system. In the output, monitor the number of data extents that are eligible to be deleted from the system. This metric has a direct correlation to the amount of free space that is available within the container storage pool.</p> • To display the amount of physical space that is occupied by a file space after the removal of the data deduplication savings, use the <code>select * from occupancy</code> command. The command output includes the <code>LOGICAL_MB</code> value. <code>LOGICAL_MB</code> is the amount of space that is used by the file space.
<p>Monitor and maintain tape devices.</p>	<p>Monitor your environment for hardware errors on tape drives and tape libraries. For instructions, see Monitoring tape alert messages for hardware errors.</p> <p>Monitor media compatibility to prevent errors on tape drives. For instructions, see Preventing errors caused by media incompatibility.</p> <p>Monitor cleaning messages for tape drives. For instructions, see Operations with cleaner cartridges.</p>	

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Evaluate the timing of client schedules. Ensure that the start and end times of client schedules do not overlap with server maintenance tasks. Limit the time for client backup operations to 8 - 12 hours.</p>	<p>On the Operations Center Overview page, click Clients > Schedules.</p> <p>In the Schedules table, the Start column displays the configured start time for the scheduled operation. To see when the most recent operation was started, hover over the clock icon.</p>	<p>Tip: You can receive a warning message if a client operation runs longer than expected. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over Clients and click Schedules. 2. Select a schedule and click Details. 3. View the details of a schedule by clicking the blue arrow next to the row. 4. In the Run time alert field, specify the time when a warning message will be issued if the scheduled operation is not completed. 5. Click Save.
<p>Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks do not overlap with client schedules.</p>	<p>On the Operations Center Overview page, click Servers > Maintenance.</p> <p>In the Maintenance table, review the information in the Last Run Time column. To see when the last maintenance task was started, hover over the clock icon.</p>	<p>The preferred method is to ensure that each maintenance task runs to completion before the next maintenance task starts. Examples of maintenance tasks include inventory expiration, copying of storage pools, space reclamation, and database backup.</p> <p>Tip: If a maintenance task is running too long, change the start time or the maximum runtime. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon and click Command Builder. 2. To change the start time or maximum runtime for a task, issue the UPDATE SCHEDULE command. For information about this command, see UPDATE SCHEDULE (Update a client schedule).

- **Monitoring tape alert messages for hardware errors**
Tape alert messages are generated by tape and library devices to report hardware errors. These messages help to determine problems that are not related to the IBM Spectrum Protect server.
- **Preventing errors caused by media incompatibility**
By monitoring and resolving media compatibility issues, you can prevent errors in an IBM Spectrum Protect tape-based solution. A new drive might have a limited ability to use media formats that are supported by a previous version of the drive. Often, a new drive can read but not write to the previous media format.
- **Operations with cleaner cartridges**
To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Related reference:

[QUERY ACTLOG](#) (Query the activity log)

Monitoring tape alert messages for hardware errors

Tape alert messages are generated by tape and library devices to report hardware errors. These messages help to determine problems that are not related to the IBM Spectrum Protect™ server.

About this task

A log page is created and can be retrieved at any time or at a specific time such as when a drive is dismantled.

A tape alert message can have one of the following severity levels:

- Informational (for example, trying to load a cartridge type that is not supported)
- Warning (for example, a hardware failure is predicted)
- Critical (for example, there is a problem with the tape and data is at risk)

Tape alert messages are turned off by default.

Procedure

- To enable tape alert messages, issue the SET TAPEALERTMSG command and specify the ON value: `set tapealertmsg on`
- To check whether tape alert messages are enabled, issue the QUERY TAPEALERTMSG command: `query tapealertmsg`

Preventing errors caused by media incompatibility

By monitoring and resolving media compatibility issues, you can prevent errors in an IBM Spectrum Protect™ tape-based solution. A new drive might have a limited ability to use media formats that are supported by a previous version of the drive. Often, a new drive can read but not write to the previous media format.

About this task

By default, existing volumes with a status of `FILLING` remain in that state after a drive upgrade. In some cases, you might want to continue to use a previous drive to fill these volumes. This preserves read/write capability for the existing volumes until they are reclaimed. If you choose to upgrade all of the drives in a library, verify that the media formats are supported by the new hardware. Unless you plan to use only the most current media with your new drive, you need to be aware of any compatibility issues. For migration instructions, see *Migrating data to upgraded drives*.

To use a new drive with media that it can read but not write to, issue the `UPDATE VOLUME` command to set the access for those volumes to read-only. This prevents errors that are caused by read/write incompatibility. For example, a new drive might eject media that is written in a format that the drive does not support as soon as the media is loaded into the drive. Or a new drive might fail the first write command to media partially written in a format that the drive does not support.

When data on the read-only media expires and the volume is reclaimed, replace it with media that is fully compatible with the new drive. Errors can be generated if a new drive is unable to correctly calibrate a volume that is written when you use a previous format. To avoid this problem, ensure that the original drive is in good working order and at current microcode levels.

Operations with cleaner cartridges

To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Monitoring the cleaning process

If a cleaner cartridge is checked in to a library, and a drive must be cleaned, the server dismounts the data volume and runs the cleaning operation. If the cleaning operation fails or is canceled, or if no cleaner cartridge is available, you might not be aware that the drive needs cleaning. Monitor cleaning messages for these problems to ensure that drives are cleaned as needed. If necessary, issue the `CLEAN DRIVE` command to have the server try the cleaning again, or manually load a cleaner cartridge into the drive.

Using multiple cleaner cartridges

The server uses a cleaner cartridge for the number of cleanings that you specify when you check in the cleaner cartridge. If you check in two or more cleaner cartridges, the server uses only one of the cartridges until the designated number of cleanings for that cartridge is reached. Then, the server uses the next cleaner cartridge. If you check in two or more cleaner cartridges and issue two or more `CLEAN DRIVE` commands concurrently, the server uses multiple cartridges at the same time and decrements the remaining cleanings on each cartridge.

Related reference:

- [AUDIT LIBRARY](#) (Audit volume inventories in an automated library)
- [CHECKIN LIBVOLUME](#) (Check a storage volume into a library)
- [CLEAN DRIVE](#) (Clean a drive)
- [LABEL LIBVOLUME](#) (Label a library volume)

Related information:

- [QUERY LIBVOLUME](#) (Query a library volume)

Verifying license compliance

Verify that your IBM Spectrum Protect™ solution complies with the provisions of your licensing agreement. By verifying compliance regularly, you can track trends in data growth or processor value unit (PVU) usage. Use this information to plan for future license purchasing.

About this task

The method that you use to verify that your solution complies with license terms varies depending on the provisions of your IBM Spectrum Protect licensing agreement.

Front-end capacity licensing

The front-end model determines license requirements based on the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems.

Back-end capacity licensing

The back-end model determines license requirements based on the terabytes of data that are stored in primary storage pools and repositories.

Tips:

- To ensure the accuracy of front-end and back-end capacity estimates, install the most recent version of the client software on each client node.
- The front-end and back-end capacity information in the Operations Center is for planning and estimation purposes.

PVU licensing

The PVU model is based on the use of PVUs by server devices.


Important: The PVU calculations that are provided by IBM Spectrum Protect are considered estimates and are not legally binding. The PVU licensing information that is reported by IBM Spectrum Protect is not considered an acceptable substitute for the IBM® License Metric Tool.


For the most recent information about licensing models, see the information about product details and licenses at the IBM Spectrum Protect product family website. If you have questions or concerns about licensing requirements, contact your IBM Spectrum Protect software provider.

Procedure

To monitor license compliance, complete the steps that correspond to the provisions of your licensing agreement.

Tip: The Operations Center provides an email report that summarizes front-end and back-end capacity usage. Reports can be sent automatically to one or more recipients regularly. To configure and manage email reports, click Reports on the Operations Center menu bar.

Option	Description
Front-end model	<p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>The front-end capacity estimate is displayed on the Front-end Usage page.</p> <p>b. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage.</p> <p>c. To estimate capacity for clients that did not report capacity usage, go to the following FTP site, which provides measuring tools and instructions:</p> <p><code>ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools</code></p> <p>To measure front-end capacity by script, complete the instructions in the most recently available licensing guide.</p> <p>d. Add the Operations Center estimate and any estimates that you obtained by using a script.</p> <p>e. Verify that the estimated capacity complies with your licensing agreement.</p>

Option	Description
Back-end model	<p>Restriction: If the source and target replication servers do not use the same policy settings, you cannot use the Operations Center to monitor back-end capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <ol style="list-style-type: none"> On the Operations Center menu bar, hover over the settings icon  and click Licensing. Click the Back-end tab. Verify that the estimated amount of data complies with your licensing agreement.
PVU model	For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model.

Tracking system status by using email reports

Set up the Operations Center to generate email reports that summarize system status. You can configure a mail server connection, change report settings, and optionally create custom reports.

Before you begin

Before you set up email reports, ensure that the following requirements are met:

- A Simple Mail Transfer Protocol (SMTP) host server is available to send and receive reports by email. The SMTP server must be configured as an open mail relay. You must also ensure that the IBM Spectrum Protect™ server that sends email messages has access to the SMTP server. If the Operations Center is installed on a separate computer, that computer does not require access to the SMTP server.
- To set up email reports, you must have system privilege for the server.
- To specify the recipients, you can enter one or more email addresses or administrator IDs. If you plan to enter an administrator ID, the ID must be registered on the hub server and must have an email address that is associated with it. To specify an email address for an administrator, use the EMAILADDRESS parameter of the UPDATE ADMIN command.

About this task

You can configure the Operations Center to send a general operations report, a license compliance report, and one or more custom reports. You create custom reports by selecting a template from a set of commonly used report templates or by entering SQL SELECT statements to query managed servers.

Procedure

To set up and manage email reports, complete the following steps:

1. On the Operations Center menu bar, click Reports.
2. If an email server connection is not yet configured, click Configure Mail Server and complete the fields. After you configure the mail server, the general operations report and license compliance report are enabled.
3. To change report settings, select a report, click Details, and update the form.
4. Optional: To add a custom report, click + Report, and complete the fields.
Tip: To immediately run and send a report, select the report and click Send.

Results

Enabled reports are sent according to the specified settings.

What to do next

The general operations report includes an attachment. To find more detailed information, expand the sections in the attachment.

If you cannot view the image in a report, you might be using an email client that converts HTML to another format. For information about restrictions, see the Operations Center online help.

Managing operations for a tape solution

Use this information to manage operations for a tape implementation for an IBM Spectrum Protect™ server.

- **Managing the Operations Center**
The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment.
- **Managing client operations**
You can resolve client errors, manage client upgrades, and decommission client nodes that are no longer required. To free storage space on the server, you can deactivate obsolete data that is stored by application clients.
- **Managing data storage**
Manage your data for efficiency and add supported devices and media to the server to store client data.
- **Managing tape devices**
Routine tape operations include preparing tape volumes for use, controlling how and when volumes are reused, and ensuring that sufficient volumes are available. You also must respond to operator requests and manage libraries, drives, disks, paths, and data movers.
- **Managing tape drives**
You can query, update, and delete tape drives. You can also clean tape drives and configure tape drive encryption and data validation.
- **Securing the IBM Spectrum Protect server**
Secure the IBM Spectrum Protect server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.
- **Stopping and starting the server**
Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.
- **Planning to upgrade the server**
When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.
- **Preparing for an outage or system update**
Prepare IBM Spectrum Protect to maintain your system in a consistent state during a planned power outage or system update.
- **Preparing for and recovering from a disaster by using DRM**
IBM Spectrum Protect provides a disaster recovery manager (DRM) function to recover your server and client data during a disaster.

Managing the Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect™ environment.

About this task

You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line. For more information about managing the Operations Center, see *Managing the Operations Center*.

Managing client operations

You can resolve client errors, manage client upgrades, and decommission client nodes that are no longer required. To free storage space on the server, you can deactivate obsolete data that is stored by application clients.

About this task

In some cases, you can resolve client errors by stopping and starting the client acceptor. If client nodes or administrator IDs are locked, you can resolve the issue by unlocking the client node or administrator ID, and then resetting the password.

For detailed instructions about identifying and resolving client errors, see *Resolving client problems*.

For instructions about adding clients, see *Protecting applications and systems*.

- **Evaluating errors in client error logs**
You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

- Stopping and restarting the client acceptor
If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.
- Resetting passwords
If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.
- Managing client upgrades
When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.
- Decommissioning a client node
If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect server, but the workstation is no longer used, you can decommission the workstation.
- Deactivating data to free storage space
In some cases, you can deactivate data that is stored on the IBM Spectrum Protect server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

Evaluating errors in client error logs

You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

Before you begin

Optionally, to resolve errors in a backup-archive client on a Linux or Windows operating system, ensure that the client management service is installed and started. For installation instructions, see [Installing the client management service](#).

Procedure

To diagnose and resolve client errors, take one of the following actions:

- If the client management service is installed on the client node, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click Details.
 3. On the client Summary page, click the Diagnosis tab.
 4. Review the retrieved log messages.

Tips:

- To show or hide the Client Logs pane, double-click the Client Logs bar.
- To resize the Client Logs pane, click and drag the Client Logs bar.

If suggestions are displayed on the Diagnosis page, select a suggestion. In the Client Logs pane, client log messages to which the suggestion relates are highlighted.

5. Use the suggestions to resolve the problems that are indicated by the error messages.
Tip: Suggestions are provided for only a subset of client messages.
- If the client management service is not installed on the client node, review the error logs for the installed client.

Stopping and restarting the client acceptor

If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.

About this task

In some cases, you can resolve client scheduling problems by stopping and restarting the client acceptor. The client acceptor must be running to ensure that scheduled operations can occur on the client. For example, if you change the IP address or domain name of the server, you must restart the client acceptor.

Procedure

Follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

- To stop the client acceptor, complete the following steps:
 - a. Determine the process ID for the client acceptor by issuing the following command on the command line:

```
ps -ef | grep dsmcad
```

Review the output. In the following sample output, 6764 is the process ID for the client acceptor:

```
root 6764 1 0 16:26:35 ? 0:00 /usr/bin/dsmcad
```

- b. Issue the following command on the command line:

```
kill -9 PID
```

where *PID* specifies the process ID for the client acceptor.

- To start the client acceptor, issue the following command on the command line:

```
/usr/bin/dsmcad
```

Linux

- To stop the client acceptor (and not restart it), issue the following command:

```
# service dsmcad stop
```

- To stop and restart the client acceptor, issue the following command:

```
# service dsmcad restart
```

MAC OS X

Click Applications > Utilities > Terminal.

- To stop the client acceptor, issue the following command:

```
/bin/launchctl unload -w com.ibm.tivoli.dsmcad
```

- To start the client acceptor, issue the following command:

```
/bin/launchctl load -w com.ibm.tivoli.dsmcad
```

Windows

- To stop the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Stop and OK.
- To restart the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Start and OK.

Related reference:

[Resolving client scheduling problems](#)

Resetting passwords

If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

Procedure

To resolve password issues, take one of the following actions:

- If a backup-archive client is installed on a client node, and the password is lost or forgotten, complete the following steps:
 1. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwreset=yes
```

where *node_name* specifies the client node and *new_password* specifies the password that you assign.

2. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to `generate` in the client options file.

- If an administrator is locked out because of password issues, complete the following steps:
 1. To provide the administrator with access to the server, issue the UNLOCK ADMIN command. For instructions, see UNLOCK ADMIN (Unlock an administrator).
 2. Set a new password by using the UPDATE ADMIN command:

```
update admin admin_name new_password forcepwreset=yes
```

where *admin_name* specifies the name of the administrator and *new_password* specifies the password that you assign.

- If a client node is locked, complete the following steps:
 1. Determine why the client node is locked and whether it must be unlocked. For example, if the client node is decommissioned, the client node is being removed from the production environment. You cannot reverse the decommission operation, and the client node remains locked. A client node also might be locked if the client data is the subject of a legal investigation.
 2. If you must unlock a client node, use the UNLOCK NODE command. For instructions, see UNLOCK NODE (Unlock a client node).
 3. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwreset=yes
```

where *node_name* specifies the name of the node and *new_password* specifies the password that you assign.

4. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to `generate` in the client options file.

Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Before you begin

1. Review the client/server compatibility requirements in technote 1053218. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted.
2. Verify system requirements for the client in IBM Spectrum Protect™ Supported Operating Systems.
3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with servers that are configured as library managers. See technote 1302789.

If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, complete the instructions that are listed in the following table.

Software	Link to instructions
IBM Spectrum Protect backup-archive client	<ul style="list-style-type: none"> • Scheduling client updates

Software	Link to instructions
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux Installing and upgrading IBM Spectrum Protect Snapshot for VMware Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> Upgrading Data Protection for SQL Server Data Protection for Oracle installation Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Virtual Environments	<ul style="list-style-type: none"> Installing and upgrading Data Protection for VMware Installing Data Protection for Microsoft Hyper-V

Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect™ server, but the workstation is no longer used, you can decommission the workstation.

About this task

When you start the decommission process, the server locks the client node to prevent it from accessing the server. Files that belong to the client node are gradually deleted, and then the client node is deleted. You can decommission the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

The simplest method for decommissioning a client node is to use the Operations Center. The decommission process runs in the background. If the client is configured to replicate client data, the Operations Center automatically removes the client from replication on the source and target replication servers before it decommissions the client.

Tip: Alternatively, you can decommission a client node by issuing the `DECOMMISSION NODE` or `DECOMMISSION VM` command. You might want to use this method in the following cases:

- To schedule the decommission process for the future or to run a series of commands by using a script, specify the decommission process to run in the background.
- To monitor the decommission process for debugging purposes, specify the decommission process to run in the foreground. If you run the process in the foreground, you must wait for the process to be completed before you continue with other tasks.

Procedure

Take one of the following actions:

- To decommission a client in the background by using the Operations Center, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click More > Decommission.
- To decommission a client node by using an administrative command, take one of the following actions:
 - To decommission an application or system client node in the background, issue the DECOMMISSION NODE command. For example, if the client node is named AUSTIN, issue the following command:


```
decommission node austin
```
 - To decommission an application or system client node in the foreground, issue the DECOMMISSION NODE command and specify the `wait=yes` parameter. For example, if the client node is named AUSTIN, issue the following command:


```
decommission node austin wait=yes
```
 - To decommission a virtual machine in the background, issue the DECOMMISSION VM command. For example, if the virtual machine is named AUSTIN, the file space is 7, and the file space name is specified by the file space ID, issue the following command:


```
decommission vm austin 7 nametype=fsid
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid
```
 - To decommission a virtual machine in the foreground, issue the DECOMMISSION VM command and specify the `wait=yes` parameter. For example, issue the following command:


```
decommission vm austin 7 nametype=fsid wait=yes
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid wait=yes
```

What to do next

Watch for error messages, which might be displayed in the user interface or in the command output, immediately after you run the process.

You can verify that the client node is decommissioned:

1. On the Operations Center Overview page, click Clients.
2. In the Clients table, in the At risk column, review the state:
 - A DECOMMISSIONED state specifies that the node is decommissioned.
 - A null value specifies that the node is not decommissioned.
 - A PENDING state specifies that the node is being decommissioned, or the decommission process failed.

Tip: If you want to determine the status of a pending decommission process, issue the following command:

```
query process
```

3. Review the command output:
 - If status is provided for the decommission process, the process is in progress. For example:

```
query process
```

Process Number	Process Description	Process Status
-----	-----	-----

- If no status is provided for the decommission process, and you did not receive an error message, the process is incomplete. A process can be incomplete if files that are associated with the node are not yet deactivated. After the files are deactivated, run the decommission process again.
- If no status is provided for the decommission process, and you receive an error message, the process failed. Run the decommission process again.

Related reference:

[DECOMMISSION NODE \(Decommission a client node\)](#)

[DECOMMISSION VM \(Decommission a virtual machine\)](#)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Spectrum Protect™ server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

About this task

Some application clients always save data to the server as active backup data. Because active backup data is not managed by inventory expiration policies, the data is not deleted automatically, and uses server storage space indefinitely. To free the storage space that is used by obsolete data, you can deactivate the data.

When you run the deactivation process, all active backup data that was stored before the specified date becomes inactive. The data is deleted as it expires and cannot be restored. The deactivation feature applies only to application clients that protect Oracle databases.

Procedure

1. From the Operations Center Overview page, click Clients.
2. In the Clients table, select one or more clients and click More > Clean Up.
Command-line method: Deactivate data by using the DEACTIVATE DATA command.

Related reference:

[DEACTIVATE DATA \(Deactivate data for a client node\)](#)

Managing data storage

Manage your data for efficiency and add supported devices and media to the server to store client data.

- **Managing inventory capacity**
Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.
- **Tuning scheduled activities**
Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.
- **Optimizing operations by enabling collocation of client files**
Collocation of client files reduces the number of volume mounts that are required when users restore, retrieve, or recall many files from a storage pool. Collocation thus reduces the amount of time that is required for these operations.

Related reference:

[Storage pool types](#)

Managing inventory capacity

Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.

Before you begin

The active and archive logs have the following characteristics:

- The active log can be a maximum size of 512 GB. For more information about sizing the active log for your system, see Planning the storage arrays.
- The archive log size is limited to the size of the file system that it is installed on. The archive log size is not maintained at a predefined size like the active log. Archive log files are automatically deleted after they are no longer needed.

As a best practice, you can optionally create an archive failover log to store archive log files when the archive log directory is full.

Check the Operations Center to determine the component of the inventory that is full. Ensure that you stop the server before you increase the size of one of the inventory components.

Procedure

- To increase the disk space for the database, complete the following steps:
 - Create one or more directories for the database on separate drives or file systems.
 - Issue the EXTEND DBSPACE command to add the directory or directories to the database. The directories must be accessible to the instance user ID of the database manager. By default, data is redistributed across all database directories and space is reclaimed.
Tips:
 - The time that is needed to complete redistribution of data and reclaiming of space is variable, depending on the size of your database. Make sure that you plan adequately.
 - Ensure that the directories that you specify are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.
 - Halt and restart the server to fully use the new directories.
 - Reorganize the database if necessary. Index and table reorganization for the server database can help to avoid unexpected database growth and performance issues. For more information about reorganizing the database, see technote 1683633.
- To decrease the size of the database for V7.1 servers and later, see the information in technote 1683633.
Restriction: The commands can increase I/O activity, and might affect server performance. To minimize performance problems, wait until one command is completed before you issue the next command. The DB2® commands can be issued when the server is running.
- To increase or decrease the size of the active log, complete the following steps:
 1. Ensure that the location for the active log has enough space for the increased log size.
 2. Halt the server.
 3. In the dsmserv.opt file, update the ACTIVELOGSIZE option to the new size of the active log, in megabytes.
The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

To change the active log to its maximum size of 512 GB, enter the following server option:

```
activelogsiz 524288
```

4. If you plan to use a new active log directory, update the directory name that is specified in the ACTIVELOGDIRECTORY server option. The new directory must be empty and must be accessible to the user ID of the database manager.
 5. Restart the server.
- Compress the archive logs to reduce the amount of space that is required for storage. Enable dynamic compression of the archive log by issuing the following command:

```
setopt archlogcompress yes
```

Restriction: Use caution when you enable the ARCHLOGCOMPRESS server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log

directory file system usage nears out of space conditions, the ARCHLOGCOMPRESS server option must be disabled. You can use the SETOPT command to disable archive log compression immediately without halting the server.

Related reference:

- [ACTIVELOGSIZE server option](#)
- [EXTEND DBSPACE \(Increase space for the database\)](#)
- [SETOPT \(Set a server option for dynamic update\)](#)

Tuning scheduled activities

Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Procedure

1. Monitor system performance regularly to ensure that backup and maintenance tasks are completing successfully. For more information about monitoring, see [Monitoring a tape solution](#).
2. If the monitoring information shows that the server workload increased, you might need to review the planning information. Review whether the capacity of the system is adequate in the following cases:
 - o The number of clients increases
 - o The amount of data that is being backed up increases
 - o The amount of time that is available for backups changes
3. Determine whether your solution has performance issues. Review the client schedules to check whether tasks are completing within the scheduled time frame:
 - a. On the Clients page of the Operations Center, select the client.
 - b. Click Details.
 - c. From the client Summary page, review the Backed up and Replicated activity to identify any risks.Adjust the time and frequency of client backup operations, if necessary.
4. Schedule enough time for the following maintenance tasks to complete successfully within a 24-hour period:
 - a. Back up the database
 - b. Run expiration to remove client backups and archive file copies from server storage.

Related concepts:

- [Performance](#)

Related tasks:

- [Deduplicating data \(V7.1.1\)](#)

Optimizing operations by enabling collocation of client files

Collocation of client files reduces the number of volume mounts that are required when users restore, retrieve, or recall many files from a storage pool. Collocation thus reduces the amount of time that is required for these operations.

About this task

With collocation enabled, the server tries to keep files on a minimal number of sequential-access storage volumes. The files can belong to a single client node, a group of client nodes, a client file space, or a group of file spaces. You can set collocation for each sequential-access storage pool when you define or update the pool.

Figure 1 shows an example of collocation by client node with three clients, each having a separate volume that contains that client's data.

Figure 1. Example of collocation enabled by node




Figure 2 shows an example of collocation by group of client nodes. Three groups are defined, and the data for each group is stored on separate volumes.

Figure 2. Example of collocation enabled by node collocation group






Figure 3 shows an example of collocation by file space group. Six groups are defined. Each group contains data from file spaces that belong to a single node. The data for each group is stored on a separate volume.

Figure 3. Example of collocation enabled by file space collocation group



When collocation is disabled, the server tries to use all available space on each volume before it selects a new volume. While this process provides better use of individual volumes, user files can become scattered across many volumes. Figure 4 shows an example of collocation that is disabled, with three clients that share space on single volume.

Figure 4. Example of collocation disabled



With collocation disabled, more media mount operations might be required to mount volumes when users restore, retrieve, or recall many files.

Collocation by group is the IBM Spectrum Protect™ system default for primary sequential-access storage pools. The default for copy storage pools is no collocation.

- Effects of collocation on operations
The effect of collocation on resources and system performance depends on the type of operation that is being run.
- Selecting volumes with collocation enabled
Volume selection depends on whether collocation is by group, node, or file space.
- Selecting volumes with collocation disabled
When collocation is disabled, the server attempts to use all available space in a storage volume before it accesses another volume.
- Collocation settings
After you define a storage pool, you can change the collocation setting by updating the storage pool. The change in collocation for the pool does not affect files that are already stored in the pool.
- Collocation of copy storage pools
Using collocation on copy storage pools requires special consideration. Collocation of copy storage pools, especially by node or file space, results in more partially filled volumes and potentially unnecessary offsite reclamation activity.
- Planning for and enabling collocation
Understanding the effects of collocation can help reduce the number of media mounts, make better use of space on sequential volumes, and improve the efficiency of server operations.

Effects of collocation on operations

The effect of collocation on resources and system performance depends on the type of operation that is being run.

Table 1 summarizes the effects of collocation on operations.

Table 1. Effect of collocation on operations

Operation	Collocation enabled	Collocation disabled
Backing up, archiving, or migrating client files	More media mounts to collocate files.	Fewer media mounts are required.
Restoring, retrieving, or recalling client files	Large numbers of files can be restored, retrieved, or recalled more quickly because files are on fewer volumes.	Multiple mounts of media might be required for a single user because files might be spread across multiple volumes. More than one user's files can be stored on the same sequential-access storage volume. For example, if two users try to recover a file that is on the same volume, the second user is forced to wait until the first user's files are recovered.
Storing data on tape	The server attempts to use all available tape volumes to separate user files before it uses all available space on every tape volume.	The server attempts to use all available space on each tape volume before the server use another tape volume.
Media mount operations	More mount operations are required when user files are backed up, archived, or migrated from client nodes directly to sequential-access volumes. More mount operations are required during reclamation and storage pool migration. More volumes are managed because volumes are not fully used.	More mount operations are required during restore, retrieve, and recall of client files.
Generating backup sets	Less time is spent searching database entries, and fewer mount operations are required.	More time is spent searching database entries and fewer mount operations are required.

When collocation is enabled for a group, single client node, or file space, all the data that belongs to the group, the node, or the file space is moved or copied by one server process. For example, if data is collocated by group, all data for all nodes that belong to the same collocation group is migrated by the same process.

When collocating data, the IBM Spectrum Protect™ server tries to keep files together on a minimal number of sequential-access storage volumes. However, when the server is backing up data to volumes in a sequential-access storage pool, the backup process has priority over collocation settings. As a result, the server completes the backup operation, but might not be able to collocate the data.

For example, suppose that you are collocating by node and you specify that a node can use two mount points on the server. Suppose also that the data that is backed up from the node can easily fit on one tape volume. During backup, the server might mount two tape volumes, and the node's data might be distributed across two tapes, rather than one. If you enable collocation, the following server operations use one server process:

- Moving data from random-access and sequential-access volumes
- Moving node data from sequential-access volumes
- Backing up a random-access or sequential-access storage pool
- Restoring a sequential-access storage pool
- Reclaiming space in a sequential-access storage pool or offsite volumes
- Migrating data from a random-access storage pool

When you migrate data from a random-access disk storage pool to a sequential-access storage pool, and collocation is by node or file space, nodes or file spaces are automatically selected for migration based on the amount of data to be migrated. The node or file space with the most data is migrated first. If collocation is by group, all nodes in the storage pool are evaluated to determine which node has the most data. The node with the most data is migrated first along with all the data for all the nodes that belong to that collocation group. This process occurs, regardless of how much data is stored in the file spaces of nodes and regardless of whether the low migration threshold was reached.

However, when you migrate collocated data from a sequential-access storage pool to another sequential-access storage pool, the server orders the volumes according to the date when the volume was last accessed. The volume with the earliest access date is migrated first, and the volume with the latest access date is migrated last.

One reason to collocate by group is that individual client nodes often do not have sufficient data to fill high-capacity tape volumes. Collocating data by groups of nodes can reduce unused tape capacity by putting more collocated data on individual tapes. Also, collocating data by groups of file spaces reduces the unused tape to a greater degree.

The data that belongs to all the nodes in the same collocation group are migrated by the same process. Therefore, collocation by group can reduce the number of times that a volume to be migrated must be mounted. Collocation by group can also minimize database scanning and reduce tape passes during data transfer from one sequential-access storage pool to another.

Selecting volumes with collocation enabled

Volume selection depends on whether collocation is by group, node, or file space.

Table 1 shows how the IBM Spectrum Protect™ server selects the first volume when collocation is enabled for a storage pool at the client-node, collocation-group, and file-space level.

Table 1. How the server selects volumes when collocation is enabled

Volume Selection Order	When collocation is by group	When collocation is by node	When collocation is by file space
1	A volume that already contains files from the collocation group to which the client belongs	A volume that already contains files from the same client node	A volume that already contains files from the same file space of that client node
2	An empty predefined volume	An empty predefined volume	An empty predefined volume
3	An empty scratch volume	An empty scratch volume	An empty scratch volume
4	A volume with the most available free space among volumes that already contain data	A volume with the most available free space among volumes that already contain data	A volume that contains data from the same client node
5	Not applicable	Not applicable	A volume with the most available free space among volumes that already contain data

When the server must continue to store data on a second volume, it uses the following selection order to acquire more space:


1. An empty predefined volume
2. An empty scratch volume
3. A volume with the most available free space among volumes that already contain data
4. Any available volume in the storage pool

When collocation is by client node or file space, the server tries to provide the best use of individual volumes and minimizes file mixing from different clients or file spaces on volumes. This configuration is depicted in Figure 1, which shows that volume selection is *horizontal*, where all available volumes are used before all available space on each volume is used. A, B, C, and D represent files from four different client nodes.

Tips:

1. If collocation is by node and the node has multiple file spaces, the server does not attempt to collocate those file spaces.
2. If collocation is by file space and a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

Figure 1. Using all available sequential-access storage volumes with collocation enabled at the node or file space level




Collocation can be by file space group or node group. When collocation is by node group (node collocation group), the server tries to collocate data from nodes that belong to the same collocation group. A file space collocation group uses the same methods as a node collocation group, but can use more space because of the granularity of file space sizes. As shown in Figure 2, data for the following groups of nodes was collocated:

- Group 1 consists of nodes A, B, and C
- Group 2 consists of nodes D and E
- Group 3 consists of nodes F, G, H, and I

Whenever possible, the IBM Spectrum Protect server collocates data that belongs to a group of nodes on a single tape, as represented by Group 2 in the figure. Data for a single node can also be spread across several tapes that are associated with a group (Group 1 and 2). If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

Figure 2. Using all available sequential-access storage volumes with collocation enabled at the group level



Normally, the IBM Spectrum Protect server always writes data to the current filling volume for the operation that is running. However, occasionally you might notice more than one filling volume in a collocated storage pool. Having more than one filling volume in a collocated storage pool can occur if different server processes or client sessions try to store data into the collocated pool at the same time. In this situation, IBM Spectrum Protect allocates a volume for each process or session that needs a volume so that both operations are completed as quickly as possible.

Selecting volumes with collocation disabled

When collocation is disabled, the server attempts to use all available space in a storage volume before it accesses another volume.


When you store client files in a sequential-access storage pool where collocation is disabled, the server selects a volume by using the following selection order:

1. A previously used sequential volume with available space (a volume with the most amount of data is selected first)
2. An empty volume

When the server needs to continue to store data on a second volume, it attempts to select an empty volume. If no empty volume exists, the server attempts to select any remaining available volume in the storage pool.

Figure 1 shows that volume use is vertical when collocation is disabled. In this example, fewer volumes are used because the server attempts to use all available space by mixing client files on individual volumes. A, B, C, and D represent files from four different client nodes.

Figure 1. Using all available space on sequential-access volumes with collocation disabled



Collocation settings

After you define a storage pool, you can change the collocation setting by updating the storage pool. The change in collocation for the pool does not affect files that are already stored in the pool.

For example, if collocation is off for a storage pool and you turn it on, from then on client files that are stored in the pool are collocated. Files that were previously stored in the storage pool are not moved to collocate them. As volumes are reclaimed, the data in the pool tends to become more collocated. You can also use the MOVE DATA or MOVE NODEDATA commands to move data to new volumes to increase collocation. Moving data to new volumes causes an increase in the processing time and the volume mount activity.

Tip: A mount wait can occur or take longer than usual when collocation by file space is enabled and a node has a volume that contains multiple file spaces. If a volume is eligible to receive data, IBM Spectrum Protect™ waits for that volume.

Collocation of copy storage pools

Using collocation on copy storage pools requires special consideration. Collocation of copy storage pools, especially by node or file space, results in more partially filled volumes and potentially unnecessary offsite reclamation activity.

Primary storage pools play a different recovery role than copy storage pools. Normally, you use primary storage pools to recover data to clients directly. In a disaster, when both clients and the server are lost, you might use offsite copy storage pool volumes to recover the primary storage pools. The types of recovery scenarios can help you to determine whether to use collocation on your copy storage pools.

Collocation typically results in partially filled volumes when you collocate by node or by file space. However, partially filled volumes are less prevalent when you collocate by group. Partially filled volumes might be acceptable for primary storage pools because the volumes remain available and can be filled during the next migration process. However, partially filled volumes might be unacceptable for copy storage pools whose storage pool volumes are taken offsite immediately. If you use collocation for copy storage pools, you must make the following decisions:

- Taking more partially filled volumes offsite, which increases the reclamation activity when the reclamation threshold is lowered or reached.
- Leaving these partially filled volumes onsite until they fill and risk not having an offsite copy of the data on these volumes.
- Whether to collocate by group to use as much tape capacity as possible.

When collocation is disabled for a copy storage pool, typically only a few partially filled volumes remain after data is backed up to the copy storage pool.

Consider your options carefully before you use collocation for copy storage pools, and whether to use simultaneous write. If you do not use simultaneous write and you use collocation for your primary storage pools, you might want to disable collocation for copy storage pools. Collocation of copy storage pools might be desirable if you have few clients with each of them having large

amounts of incremental backup data each day. For collocation with simultaneous write, you must ensure that the collocation settings are identical for the primary storage pools and copy storage pools.

Planning for and enabling collocation

Understanding the effects of collocation can help reduce the number of media mounts, make better use of space on sequential volumes, and improve the efficiency of server operations.

About this task

Table 1 lists the four collocation options that you can specify on the DEFINE STGPOOL and UPDATE STGPOOL commands. The table also shows the effects of collocation on data that belongs to nodes that are and are not members of collocation groups.

Table 1. Collocation options and the effects on node data

Collocation option	If a node is not defined as a member of a collocation group	If a node is defined as a member of a collocation group
No	The data for the node is not collocated.	The data for the node is not collocated.
Group	The server stores the data for the node on as few volumes in the storage pool as possible.	The server stores the data for the node and for other nodes that belong to the same collocation group on as few volumes as possible.
Node	The server stores the data for the node on as few volumes as possible.	The server stores the data for the node on as few volumes as possible.
File space	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.

Table 2. Collocation group options and effects on file space data

Collocation option	If a file space is not defined as a member of a collocation group	If a file space is defined as a member of a collocation group
No	The data for the file space is not collocated.	The data for the file space is not collocated.
Group	The server stores the data for the file space on as few volumes in the storage pool as possible.	The server stores the data for the file space and other file spaces that belong to the same collocation group on as few volumes as possible.
Node	The server stores the data for the node on as few volumes as possible.	The server stores the data for the node on as few volumes as possible.
File space	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.	The server stores the data for the file spaces on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.

Procedure

To determine whether and how to collocate data, complete the following steps:

1. Determine how to organize data, whether by client node, group of client nodes, or file space. To collocate by group, you must decide how to group nodes:
 - o If the goal is to save space, you might want to group small nodes together to better use tapes.
 - o If the goal is potentially faster client restores, group nodes together so that they fill as many tapes as possible. By grouping nodes together, the individual node data is distributed across two or more tapes and that more tapes can be mounted simultaneously during a multi-session no-query restore operation.
 - o If the goal is to departmentalize data, you can group nodes by department.
2. To collocate groups, complete the following steps:
 - a. Define collocation groups with the DEFINE COLLOGROUP command.
 - b. Add client nodes to the collocation groups with the DEFINE COLLOCMEMBER command.

The following query commands are available to help in collocating groups:

QUERY COLLOGROUP

Displays the collocation groups that are defined on the server.

QUERY NODE

Displays the collocation group, if any, to which a node belongs.

QUERY NODEDATA

Displays information about the data for one or more nodes in a sequential-access storage pool.

QUERY STGPOOL

Displays information about the location of client data in a sequential-access storage pool and the amount of space a node occupies in a volume.

You can also use IBM Spectrum Protect™ server scripts or Perl scripts to display information that can be useful in defining collocation groups.

3. Specify how data must be collocated in a storage pool by issuing the DEFINE STGPOOL or UPDATE STGPOOL command and specifying the COLLOCATE parameter.

What to do next

Tip: To reduce the number of media mounts, use space on sequential volumes more efficiently, and enable collocation, complete the following steps:

- Define a storage pool hierarchy and policy to require that backed-up, archived, or space-managed files are initially stored in disk storage pools.

When files are migrated from a disk storage pool, the server attempts to migrate all files that belong to the client node or collocation group that is using the most disk space in the storage pool. This process works well with the collocation option because the server tries to place all of the files from a particular client on the same sequential-access storage volume.

- Use scratch volumes for sequential-access storage pools to allow the server to select new volumes for collocation.
- Specify the client option COLLOCATEBYFILESPEC to limit the number of tapes to which objects associated with one file specification are written. This collocation option makes collocation by the server more efficient; it does not override collocation by file space or collocation by node.

Managing tape devices

Routine tape operations include preparing tape volumes for use, controlling how and when volumes are reused, and ensuring that sufficient volumes are available. You also must respond to operator requests and manage libraries, drives, disks, paths, and data movers.

- **Preparing removable media**
You must prepare removable media before it can be used to store data. Typical preparation tasks include labeling and checking in volumes.
- **Managing volume inventory**
You can manage volume inventory by controlling the server's access to volumes, by reusing tapes, and by reusing volumes that are used for database backup and export operations. You can also manage inventory by maintaining a supply of scratch volumes.
- **Partially written volumes**
Partially written volumes are always private volumes, even if their status was scratch before the server mounted them. The server tracks the original status of scratch volumes and returns them to scratch status when they are empty.
- **Operations with shared libraries**
Shared libraries are logical libraries that are represented physically by SCSI libraries. The physical library is controlled by the IBM Spectrum Protect server that is configured as a library manager. IBM Spectrum Protect servers that use the SHARED library type are library clients to the IBM Spectrum Protect library manager server.
- **Managing server requests for volumes**
IBM Spectrum Protect displays requests and status messages to all administrative command-line clients that are started in console mode. These request messages often have a time limit. Successful server operations must be completed within the time limit that is specified; otherwise, the operation times out.

Preparing removable media

You must prepare removable media before it can be used to store data. Typical preparation tasks include labeling and checking in volumes.

About this task

When IBM Spectrum Protect™ accesses a removable media volume, it verifies the volume name in the label header to ensure that the correct volume is accessed.

Tape volumes must be labeled before the server can use them.

Procedure

To prepare a volume for use, complete the following steps:

1. Label the volume by issuing the LABEL LIBVOLUME command.
2. For automated libraries, check the volume into the library. For instructions, see [Checking volumes into an automated library](#),
Tip: When you use the LABEL LIBVOLUME command with drives in an automated library, you can label and check in the volumes with one command.
3. If the storage pool cannot contain scratch volumes (MAXSCRATCH=0), identify the volume to IBM Spectrum Protect by name so that the volume can be accessed later.

If the storage pool can contain scratch volumes (MAXSCRATCH is set to a non-zero value), skip this step.

- Labeling tape volumes
You must label tape volumes before the server can use them.
- Checking volumes into an automated library
You can check in a volume to an automated library by using the CHECKIN LIBVOLUME command.

Labeling tape volumes

You must label tape volumes before the server can use them.

About this task

For automated libraries, you are prompted to insert the volume in the entry/exit slot of the library. If no convenience input/output (I/O) station is available, insert the volume into an empty slot. You can label the volumes when you check them in or before you check them in.

Procedure

To label tape volumes before you check them in, complete the following steps:

1. Label tape volumes by issuing the LABEL LIBVOLUME command. For example, to name a library volume VOLUME1 in a library that is named LIBRARY 1, issue the following command:

```
label libvolume library1 volume1
```

Requirement: At least one drive must be available. The drive cannot be used by another IBM Spectrum Protect™ process. If a drive is idle, the drive is considered to be unavailable.

2. To overwrite an existing label, specify the OVERWRITE=YES parameter. By default, the LABEL LIBVOLUME command does not overwrite an existing label.
- Labeling volumes in a SCSI library
You can label volumes individually or use IBM Spectrum Protect to search the library for volumes and label the found volumes.

Related tasks:

[Labeling new volumes by using AUTOLABEL](#)

Related reference:

[LABEL LIBVOLUME \(Label a library volume\)](#)

Checking volumes into an automated library

You can check in a volume to an automated library by using the CHECKIN LIBVOLUME command.

Before you begin

To automatically label tapes before you check them in, issue the `DEFINE LIBRARY` command and specify the `AUTOLABEL=YES` parameter. By using the `AUTOLABEL` parameter, you eliminate the need to prelabel a set of tapes.

About this task

Each volume that is used by a server for any purpose must have a unique name. This requirement applies to all volumes, whether the volumes are used for storage pools, or used for operations such as database backup or export. The requirement also applies to volumes that are in different libraries but that are used by the same server.

Tips:

- Do not use a single library for volumes that have bar code labels and volumes that do not have bar code labels. Bar code scanning can take a long time for unlabeled volumes.
- The server accepts only tapes that are labeled with IBM® standard labels.
- Any volume that has a bar code that begins with `CLN` is treated as a cleaning tape.
- If a volume has an entry in volume history, you cannot check it in as a scratch volume.

Procedure

1. To check a storage volume into a library, issue the `CHECKIN LIBVOLUME` command.
Tip: The command always runs as a background process. Wait for the `CHECKIN LIBVOLUME` process to complete processing before you define volumes, or the defining process fails. You can save time by checking in volumes as part of the labeling operation.
 2. Name the library and specify whether the volume is a private volume or a scratch volume. Depending on whether you use scratch volumes or private volumes, complete one of the following steps:
 - If you use only scratch volumes, ensure that enough scratch volumes are available. For example, you might need to label more volumes. As volumes are used, you might also need to increase the number of scratch volumes that are allowed in the storage pool that you defined for this library.
 - If you want to use private volumes in addition to or instead of scratch volumes in the library, define volumes to the storage pool by using the `DEFINE VOLUME` command. You must label and check in the volumes that you define.
- **Checking a single volume into a SCSI library**
You can check in a single volume by issuing the `CHECKIN LIBVOLUME` command and specifying the `SEARCH=NO` parameter. IBM Spectrum Protect™ requests that the mount operator load the volume into the entry/exit port of the library.
 - **Checking in volumes from library storage slots**
When you have many volumes to check in and you want to avoid issuing a `CHECKIN LIBVOLUME` command for each volume, you can search storage slots for new volumes. The server finds volumes that have not yet been added to the volume inventory.
 - **Checking in volumes from library entry/exit ports**
You can search all slots of bulk entry/exit ports for labeled volumes and the server can check them in automatically.
 - **Checking in volumes by using library bar code readers**
You can save time when you check in volumes to libraries that have bar code readers by using the characters on the bar code labels as names for the volumes.
 - **Checking in volumes by using a bar code reader**
You can save time when you check in volumes by using a bar code reader, if your library has one.
 - **Checking volumes into a full library with swapping**
If no empty slots are available in the library when you are checking in volumes, the check-in operation fails unless you enable *swapping*. If you enable swapping and the library is full, the server selects a volume to eject and then checks in the volume that you requested.
 - **Windows** Private volumes and scratch volumes
To optimize tape storage, review the information about private volumes and scratch volumes. Use private volumes and scratch volumes appropriately.
 - **Windows** Element addresses for library storage slots
An element address is a number that indicates the physical location of a storage slot or drive within an automated library.

Related tasks:

Labeling tape volumes

Checking a single volume into a SCSI library

You can check in a single volume by issuing the CHECKIN LIBVOLUME command and specifying the SEARCH=NO parameter. IBM Spectrum Protect™ requests that the mount operator load the volume into the entry/exit port of the library.

Procedure

1. Issue the CHECKIN LIBVOLUME command.

For example, to check in volume VOL001, enter the following command:

```
checkin libvolume tapelib vol001 search=no status=scratch
```

2. Respond to the prompt from the server.
 - o If the library has an entry/exit port, you are prompted to insert a tape into the entry/exit port.
 - o If the library does not have an entry/exit port, you are prompted to insert a tape into one of the slots in the library. Element addresses identify these slots. For example, the server finds that the first empty slot is at element address 5. The following message is returned:

```
ANR8306I 001: Insert 8MM volume VOL001 R/W in slot with element  
address 5 of library TAPELIB within 60 minutes; issue 'REPLY' along  
with the request ID when ready.
```

If you do not know the location of element address 5 in the library, check the worksheet for the device. To find the worksheet, review the documentation for your library. After you insert the volume as requested, respond to the message from an IBM Spectrum Protect administrative client. Issue the REPLY command, followed by the request number (the number at the beginning of the mount request) for example:

```
reply 1
```

Tip: Element addresses are sometimes numbered starting with a number other than 1. Check the worksheet to be sure. If no worksheet is listed for your device in IBM® Support Portal for IBM Spectrum Protect, see the documentation for your library.

If you specify a wait time of 0 by using the optional WAITTIME parameter on the CHECKIN LIBVOLUME command, a REPLY command is not required. The default wait time is 60 minutes.

Checking in volumes from library storage slots

When you have many volumes to check in and you want to avoid issuing a CHECKIN LIBVOLUME command for each volume, you can search storage slots for new volumes. The server finds volumes that have not yet been added to the volume inventory.

Procedure

1. Open the library and place the new volumes in unused slots. For example, for a SCSI device, open the library access door, place all of the new volumes in unused slots, and close the door.
2. If the volumes are not labeled, use the LABEL LIBVOLUME command to label the volume.
3. Issue the CHECKIN LIBVOLUME command with the SEARCH=YES parameter.

Related reference:

[☞ CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Checking in volumes from library entry/exit ports

You can search all slots of bulk entry/exit ports for labeled volumes and the server can check them in automatically.

Before you begin

Issue the LABEL LIBVOLUME command to label volumes that are not labeled.

About this task

For SCSI libraries, the server scans all of the entry/exit ports in the library for volumes. If a volume is found that contains a valid volume label, it is checked in automatically.

Procedure

Issue the CHECKIN LIBVOLUME command and specify the SEARCH=BULK parameter.

- To load a tape in a drive and read the label, specify the CHECKLABEL=YES parameter. After the server reads the label, the server moves the tape from the drive to a storage slot.
- To have the server use the bar code reader to verify external labels on tapes, specify the CHECKLABEL=BARCODE parameter. When bar code reading is enabled, the server reads the label and moves the tape from the entry/exit port to a storage slot.

Checking in volumes by using library bar code readers

You can save time when you check in volumes to libraries that have bar code readers by using the characters on the bar code labels as names for the volumes.

About this task

The server reads the bar code labels and uses the information to write the internal media labels. For volumes that have no bar code labels, the server mounts the volumes in a drive and attempts to read the internal, recorded label.

Procedure

Issue the CHECKIN LIBVOLUME command with the CHECKLABEL=BARCODE parameter. For example, to use a bar code reader to search a library that is named TAPELIB and check in a scratch tape, issue the following command:

```
checkin libvolume tapelib search=yes status=scratch checklabel=barcode
```

Checking in volumes by using a bar code reader

You can save time when you check in volumes by using a bar code reader, if your library has one.

About this task

When you check in a volume, you can specify whether the media labels are read during check-in processing. When label-checking is on, IBM Spectrum Protect™ mounts each volume to read the internal label and checks in a volume only if it is correctly labeled. Label-checking can prevent future errors when volumes are used in storage pools, but also increases processing time at check-in.

If a volume has no bar code label, IBM Spectrum Protect mounts the volumes in a drive and attempts to read the recorded label.

Procedure

To check in volumes by using a bar code reader, issue the CHECKIN LIBVOLUME command and specify CHECKLABEL=BARCODE. For example, to use the bar code reader to check in all volumes as scratch volumes in a library that is named TAPELIB, issue the following command:

```
checkin libvolume tapelib search=yes status=scratch checklabel=barcode
```

Related tasks:

Preparing removable media

Related reference:

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Checking volumes into a full library with swapping

If no empty slots are available in the library when you are checking in volumes, the check-in operation fails unless you enable *swapping*. If you enable swapping and the library is full, the server selects a volume to eject and then checks in the volume that you requested.

About this task

The server selects the volume to eject by checking first for any available scratch volume, then for the volume that is least frequently mounted. The server ejects the volume that it selects for the swap operation from the library and replaces the ejected volume with the volume that is being checked in.

Procedure

To swap volumes if an empty library slot is not available to check in a volume, issue the CHECKIN LIBVOLUME command and specify the SWAP=YES parameter. For example, to check in a volume that is named VOL1 into a library that is named AUTO and specify swapping, issue the following command:

```
checkin libvolume auto voll swap=yes
```

Related tasks:

Managing a full library with an overflow location

Related reference:

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Private volumes and scratch volumes

To optimize tape storage, review the information about private volumes and scratch volumes. Use private volumes and scratch volumes appropriately.

Private volumes cannot be overwritten when a scratch mount is requested. You cannot check in a volume with scratch status when that volume is used by a storage pool, to export data, to back up a database or to back up to a backup set volume.

Partially written volumes are always private volumes. Volumes have a status of either scratch or private, but when IBM Spectrum Protect™ stores data on them, their status becomes private.

Table 1. Private volume and scratch volume uses

Type of volume	When to use
Private volumes	Use private volumes to regulate the volumes that are used by individual storage pools, and to manually control the volumes. To define private volumes, issue the DEFINE VOLUME command. For database restore, memory dumps, or loads, or for server import operations, you must specify private volumes.
Scratch volumes	In some cases, you can simplify volume management by using scratch volumes. You can use scratch volumes in the following circumstances: <ul style="list-style-type: none"> • When you do not need to define each storage pool volume. • When you want to take advantage of the automation of robotic devices. • When different storage pools share an automated library, and the storage pools can dynamically acquire volumes from the scratch volumes in the library. The volumes do not have to be preallocated to the storage pools.

Related tasks:

Changing the status of a volume in an automated library

Related reference:

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DELETE VOLUME \(Delete a storage pool volume\)](#)

Element addresses for library storage slots

An element address is a number that indicates the physical location of a storage slot or drive within an automated library.

If a library has entry/exit ports, you can add and remove media by using the ports. If no entry/exit port exists, you must load tapes into storage slots.

If you load tapes into storage slots, you must reply to mount requests that identify storage slots with element addresses. If you specify a wait time of 0 on the CHECKIN LIBVOLUME command or the LABEL LIBVOLUME command, you do not need to reply to a mount request.

For element addresses, see the device manufacturer's documentation or go to the IBM® Support Portal for IBM Spectrum Protect™ and search for element addresses.

Related reference:

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

Managing volume inventory

You can manage volume inventory by controlling the server's access to volumes, by reusing tapes, and by reusing volumes that are used for database backup and export operations. You can also manage inventory by maintaining a supply of scratch volumes.

About this task

Each volume that is used by a server must have a unique name, whether the volumes are used for storage pools, or used for operations such as database backup or export. Volumes that are in different libraries but that are used by the same server must also have a unique name.

- Controlling access to volumes
You can use different methods to control access to volumes.
- Reusing tapes
To ensure an adequate supply of tapes, you can expire old files, reclaim volumes, and delete volumes that reach end of life. You can also maintain a supply of scratch volumes.
- Maintaining a supply of scratch volumes
You must set the maximum number of scratch volumes for a storage pool high enough for the expected usage.
- **AIX** | **Linux** Maintaining a supply of volumes in a library that contains WORM media
For libraries that contain Write Once Read Many (WORM) media, you can prevent cancellation of data storage transactions by maintaining a supply of scratch or new private volumes in the library. Canceled transactions can cause WORM media to be wasted.
- Manage the volume inventory in automated libraries
The IBM Spectrum Protect™ server uses a library volume inventory to track scratch and private volumes that are available in an automated library. You must ensure that the inventory is consistent with the volumes that are physically in the library.

Controlling access to volumes

You can use different methods to control access to volumes.

Procedure

To control access to volumes, take any of the following actions:

- To prevent the server from mounting a volume, issue the UPDATE VOLUME command and specify the ACCESS=UNAVAILABLE parameter.
- To make volumes unavailable and send them offsite for protection, use a copy storage pool or an active-data storage pool.
- You can back up primary storage pools to a copy storage pool and then send the copy storage pool volumes offsite.
- You can copy active versions of client backup data to active-data storage pools, and then send the volumes offsite.
- You can track copy storage pool volumes and active-data pool volumes by changing their access mode to offsite, and updating the volume history to identify their location.

Related reference:

[UPDATE VOLUME \(Update a storage pool volume\)](#)

Reusing tapes

To ensure an adequate supply of tapes, you can expire old files, reclaim volumes, and delete volumes that reach end of life. You can also maintain a supply of scratch volumes.

About this task

Over time, media age, and you might not need some of the backup data that is stored on the media. You can define server policies to determine how many backup versions are retained and how long they are retained. You can use expiration processing to delete files that you no longer require. You can keep the data that you require on the media. When you no longer require the data, you can then reclaim and reuse the media.

Procedure

1. Delete unnecessary client data by regularly running expiration processing. Expiration processing deletes data that is no longer valid either because it exceeds the retention specifications in the policy or because users or administrators deleted the active versions of the data.
2. Reuse volumes in storage pools by running reclamation processing.

Reclamation processing consolidates any unexpired data by moving it from multiple volumes onto fewer volumes. The media can then be returned to the storage pool and reused.

3. Reuse volumes that contain outdated database backups or exported data that is no longer required by deleting volume history.

Before the server can reuse volumes that are tracked in the volume history, you must delete the volume information from the volume history file by issuing the `DELETE VOLHISTORY` command.

Tip: If your server uses the disaster recovery manager (DRM) function, the volume information is automatically deleted during `MOVE DRMEDIA` command processing.

4. Determine when tape volumes reach end of life. You can use the server to display statistics about volumes, including the number of write operations that are completed on the media and the number of write errors. Private volumes and scratch volumes display the following statistical data:

Private volumes

For media initially defined as private volumes, the server maintains this statistical data, even as the volume is reclaimed. You can compare the information with the number of write operations and write errors that are recommended by the manufacturer.

Scratch volumes

For media initially defined as scratch volumes, the server overwrites this statistical data each time the volumes are reclaimed.

5. Reclaim any valid data from volumes that reach end of life. If the volumes are in automated libraries, check them out of the volume inventory. Delete private volumes from the database with the `DELETE VOLUME` command.
6. Ensure that volumes are available for tape rotation so that the storage pool does not run out of space. You can use the Operations Center to monitor the availability of scratch volumes. Ensure that the number of scratch volumes is high enough to meet demand. For more information, see [Maintaining a supply of volumes in a library that contains WORM media](#). WORM media: Write Once Read Many (WORM) drives can waste media when the server cancels transactions because volumes are unavailable to complete the backup operation. After the server writes to WORM volumes, the space on the volumes cannot be reused, even if the transactions are canceled (for example, if a backup is canceled because of a shortage of media in the device). To minimize wasted WORM media, complete the following actions:
 - a. Ensure that the maximum number of scratch volumes for the device storage pool is at least equal to the number of storage slots in the library.
 - b. Check enough volumes into the device's volume inventory for the expected load.

If most backup operations are for small files, controlling the transaction size can affect how WORM platters are used. Smaller transactions mean that less space is wasted when a transaction such as a backup operation must be canceled. Transaction size is controlled by a server option, `TXNGROUPMAX`, and a client option, `TXNBYTELIMIT`.

Related tasks:

[Migrating data to upgraded drives](#)

[Managing server requests for volumes](#)

Related reference:

[DELETED VOLHISTORY](#) (Delete sequential volume history information)

[DELETE VOLUME](#) (Delete a storage pool volume)

- 🔗 Txnbytelimit option
- 🔗 TXNGROUPMAX server option

Related information:

- 🔗 EXPIRE INVENTORY (Manually start inventory expiration processing)
- 🔗 RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)

Maintaining a supply of scratch volumes

You must set the maximum number of scratch volumes for a storage pool high enough for the expected usage.

About this task

When you define a storage pool, you must specify the maximum number of scratch volumes that the storage pool can use. The server automatically requests a scratch volume when needed. When the number of scratch volumes that the server is using for the storage pool exceeds the specified maximum, the storage pool can run out of space.

Procedure

When a storage pool needs more than the maximum number of scratch volumes, you can take one or both of the following actions:

1. Increase the maximum number of scratch volumes by issuing the UPDATE STGPOOL command and specifying the MAXSCRATCH parameter.
2. Make volumes available for reuse by running expiration processing and reclamation to consolidate data onto fewer volumes.
 - a. Issue the EXPIRE INVENTORY command to run expiration processing.

Tip: By default this process automatically runs every day. You can also specify the EXPINTERVAL server option in the server options file, dsmserv.opt, to run expiration processing automatically. A value of 0 means that you must use the EXPIRE INVENTORY command to run expiration processing.
 - b. Issue the RECLAIM STGPOOL command to run reclamation processing.

Tip: You can also specify reclamation thresholds when you define the storage pool by using the DEFINE STGPOOL command and specifying the RECLAIMPROCESS parameter.

What to do next

If you need more volumes for future backup operations, label more scratch volumes by using the LABEL LIBVOLUME command.

Related tasks:

Maintaining a supply of scratch volumes in an automated library

Related reference:

- 🔗 LABEL LIBVOLUME (Label a library volume)
- 🔗 UPDATE STGPOOL (Update a storage pool)

Related information:

- 🔗 EXPIRE INVENTORY (Manually start inventory expiration processing)
- 🔗 RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)

Maintaining a supply of volumes in a library that contains WORM media

For libraries that contain Write Once Read Many (WORM) media, you can prevent cancellation of data storage transactions by maintaining a supply of scratch or new private volumes in the library. Canceled transactions can cause WORM media to be wasted.

About this task

IBM Spectrum Protect™ cancels a transaction if volumes, either private or scratch, are unavailable to complete the data storage operation. After IBM Spectrum Protect begins a transaction by writing to a WORM volume, the written space on the volume cannot be reused, even if the transaction is canceled.

For example, if you have WORM volumes that hold 2.6 GB each and a client starts to back up a 12 GB file. If IBM Spectrum Protect cannot acquire a fifth scratch volume after four volumes are full, IBM Spectrum Protect cancels the backup operation. The four volumes that IBM Spectrum Protect already filled cannot be reused.

To minimize cancellation of transactions, you must have enough volumes available in the library to manage expected client operations such as backups.

Procedure

1. Ensure that the storage pool that is associated with the library has sufficient scratch volumes. Issue the UPDATE STGPOOL command and specify the MAXSCRATCH parameter.
2. To manage the expected load, check in a sufficient number of scratch or private volumes to the library by issuing the CHECKIN LIBVOLUME command.
3. To control transaction size, specify the TXNGROUPMAX server option and the TXNBYTELIMIT client option. If your clients tend to store small files, controlling the transaction size can affect how WORM volumes are used. Smaller transactions waste less space when a transaction such as a backup must be canceled.

Related reference:

- [CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)
- [UPDATE STGPOOL \(Update a storage pool\)](#)
- [Txnbytelimit option](#)
- [TXNGROUPMAX server option](#)

Manage the volume inventory in automated libraries

The IBM Spectrum Protect™ server uses a library volume inventory to track scratch and private volumes that are available in an automated library. You must ensure that the inventory is consistent with the volumes that are physically in the library.

The library volume inventory is separate from the inventory of volumes for each storage pool. To add a volume to a library volume inventory, you check in a volume to that IBM Spectrum Protect library.

A list of volumes in the library volume inventory might not be identical to a list of volumes in the storage pool inventory for the device. For example, you can check in scratch volumes to the library but you cannot define them to a storage pool. If scratch volumes are not selected for backup operations, you can define private volumes to a storage pool but you cannot check them into the volume inventory for the device.

To ensure that the volume inventory for the server library remains accurate, check out volumes to physically remove the volumes from a SCSI library. When you check out a volume that is used by a storage pool, the volume remains in the storage pool. If you must mount the volume when it is checked out, a message to the mount operator's console is displayed with a request to check in the volume. If the check-in operation is unsuccessful, the server marks the volume as unavailable.

When a volume is in the library volume inventory, you can change the status of the volume from scratch to private.

To check whether the volume inventory for the server library is consistent with the volumes that are physically in the library, you can audit the library. The inventory can become inaccurate if volumes are moved in and out of the library without informing the server by using volume check-in or check-out operations.

- **Changing the status of a volume in an automated library**
You can change the status of a volume from private to scratch or from scratch to private.
- **Removing volumes from an automated library**
You can remove volumes from an automated library if you exported data to a volume and want to import the data to another system. You might also want to remove volumes to create space for new volumes.
- **Maintaining a supply of scratch volumes in an automated library**
When you define a storage pool that is associated with an automated library, you can specify a maximum number of scratch volumes equal to the physical capacity of the library. If the server is using a greater number of scratch volumes for the storage pool, you must ensure that enough volumes are available.
- **Managing a full library with an overflow location**
As the demand for storage grows, the number of volumes that you need for a storage pool might exceed the physical capacity of an automated library. To make space available for new volumes and to monitor existing volumes, you can define an overflow location for a storage pool.
- **Auditing the volume inventory in a library**
You can audit an automated library to ensure that the library volume inventory is consistent with the volumes that are physically in the library. You might want to audit a library if the library volume inventory is distorted due to manual movement of volumes in the library or to database problems.

Related tasks:

[Checking volumes into an automated library](#)

Related reference:

[AUDIT LIBRARY](#) (Audit volume inventories in an automated library)

Changing the status of a volume in an automated library

You can change the status of a volume from private to scratch or from scratch to private.

Procedure

To change the status of a volume, issue the UPDATE LIBVOLUME command. For example, to change the status of a volume that is named VOL1 to a private volume, issue the following command:

```
update libvolume lib1 vol1 status=private
```

Restrictions:

- You cannot change the status of a volume from private to scratch if the volume belongs to a storage pool or is defined in the volume history file.
- Private volumes must be administrator-defined volumes with either no data or invalid data. They cannot be partially written volumes that contain active data. Volume statistics are lost when volume statuses are modified.

Removing volumes from an automated library

You can remove volumes from an automated library if you exported data to a volume and want to import the data to another system. You might also want to remove volumes to create space for new volumes.

About this task

By default, the server mounts the volume that you check out and verifies the internal label. When the label is verified, the server removes the volume from the library volume inventory, and then moves it to the entry/exit port or convenience I/O station of the library. If the library does not have an entry/exit port, the server requests that the mount operator remove the volume from a slot or device within the library.

Procedure

- To remove a volume from an automated library, issue the CHECKOUT LIBVOLUME command.
- For automated libraries with multiple entry/exit ports, issue the CHECKOUT LIBVOLUME command and specify the REMOVE=BULK parameter. The server ejects the volume to the next available entry/exit port.

What to do next

If you check out a volume that is defined in a storage pool and the server must access the volume later, the server requests that the volume be checked in. To return volumes to a library, issue the CHECKIN LIBVOLUME command.

Related reference:

[CHECKIN LIBVOLUME](#) (Check a storage volume into a library)

[CHECKOUT LIBVOLUME](#) (Check a storage volume out of a library)

Maintaining a supply of scratch volumes in an automated library

When you define a storage pool that is associated with an automated library, you can specify a maximum number of scratch volumes equal to the physical capacity of the library. If the server is using a greater number of scratch volumes for the storage pool, you must ensure that enough volumes are available.

Procedure

If the number of scratch volumes that the server is using for the storage pool exceeds the number that is specified in the storage pool definition, complete the following steps:

1. Add scratch volumes to the library by issuing the CHECKIN LIBVOLUME command.

Tip: You might have to use an overflow location to move volumes out of the library to make room for these scratch volumes. For more information, see [Managing a full library with an overflow location](#).

2. Increase the maximum number of scratch volumes that can be added to a storage pool by issuing the UPDATE STGPOOL command and specifying the MAXSCRATCH parameter.

What to do next

You might need more volumes for future recovery operations, so consider labeling and setting aside extra scratch volumes.

Related tasks:

Maintaining a supply of scratch volumes

Managing a full library with an overflow location

As the demand for storage grows, the number of volumes that you need for a storage pool might exceed the physical capacity of an automated library. To make space available for new volumes and to monitor existing volumes, you can define an overflow location for a storage pool.

About this task

The server tracks the volumes that are moved to the overflow area and makes storage slots available for new volumes.

Procedure

1. Create a volume overflow location. Define or update the storage pool that is associated with the automated library by issuing the DEFINE STGPOOL or UPDATE STGPOOL command and specifying the OVFLOCATION parameter. For example, to create an overflow location that is named ROOM2948 for a storage pool that is named ARCHIVEPOOL, issue the following command:

```
update stgpool archivepool ovflocation=Room2948
```

2. When you need to create space in the library for scratch volumes, move full volumes to the overflow location by issuing the MOVE MEDIA command. For example, to move all full volumes in the specified storage pool out of the library, issue the following command:

```
move media * stgpool=archivepool
```

3. Check in scratch volumes as needed.
Restriction: If a volume has an entry in the volume history file, you cannot check it in as a scratch volume. For more information, see [Checking volumes into an automated library](#).
4. Identify the empty scratch tapes in the overflow location by issuing the QUERY MEDIA command. For example, issue the following command:

```
query media * stg=* whereovflocation=Room2948 wherestatus=empty
```

5. If the server requests additional volumes, locate and check in volumes from the overflow location.

To find volumes in an overflow location, issue the QUERY MEDIA command. You can also use the QUERY MEDIA command to generate commands by checking in volumes.

For example, to list the volumes in the overflow location, and at the same time generate the commands to check those volumes into the library, issue a command that is similar to the following example:

```
query media format=cmd stgpool=archivepool whereovflocation=Room2948  
cmd="checkin libvol autolib &vol status=private"  
cmdfilename="\storage\move\media\checkin.vols"
```

Tips:

- o Mount requests from the server include the location of the volumes.
- o To specify the number of days that must elapse before the volumes are eligible for processing, issue the UPDATE STGPOOL command and specify the REUSEDELAY parameter.
- o The file that contains the generated commands can be run by using the IBM Spectrum Protect™ MACRO command.

Related reference:

[MOVE MEDIA \(Move sequential-access storage pool media\)](#)

[QUERY MEDIA](#) (Query sequential-access storage pool media)

[UPDATE STGPOOL](#) (Update a storage pool)

Auditing the volume inventory in a library

You can audit an automated library to ensure that the library volume inventory is consistent with the volumes that are physically in the library. You might want to audit a library if the library volume inventory is distorted due to manual movement of volumes in the library or to database problems.

Procedure

1. Ensure that no volumes are mounted in the library drives. If any volumes are mounted in the IDLE state, issue the DISMOUNT VOLUME command to dismount them.
2. Audit the volume inventory by issuing the AUDIT LIBRARY command. Take one of the following actions:
 - o If the library has a bar code reader, you can save time by using the bar code reader to identify volumes. For example, to audit the TAPELIB library by using its bar code reader, issue the following command:

```
audit library tapelib checklabel=barcode
```

- o If the library does not have a bar code reader, issue the AUDIT LIBRARY command without specifying CHECKLABEL=BARCODE. The server mounts each volume to verify the label. After the label is verified, the server completes auditing any remaining volumes.

Results

The server deletes missing volumes from the inventory and updates the locations of volumes that moved since the last audit.

Restriction: The server cannot add new volumes to the inventory during an audit operation.

Related tasks:

Labeling tape volumes

Related reference:

[AUDIT LIBRARY](#) (Audit volume inventories in an automated library)

[DISMOUNT VOLUME](#) (Dismount a volume by volume name)

Partially written volumes

Partially written volumes are always private volumes, even if their status was scratch before the server mounted them. The server tracks the original status of scratch volumes and returns them to scratch status when they are empty.

Except for volumes in automated libraries, the server is unaware of a scratch volume until after the volume is mounted. Then, the volume status changes to private, and the volume is automatically defined as part of the storage pool for which the mount request was made.

Related tasks:

Changing the status of a volume in an automated library

Operations with shared libraries

Shared libraries are logical libraries that are represented physically by SCSI libraries. The physical library is controlled by the IBM Spectrum Protect™ server that is configured as a library manager. IBM Spectrum Protect servers that use the SHARED library type are library clients to the IBM Spectrum Protect library manager server.

The library client contacts the library manager when the library manager starts and the storage device initializes, or after a library manager is defined to a library client. The library client confirms that the contacted server is the library manager for the named library device. The library client also compares drive definitions with the library manager for consistency. The library client contacts the library manager for each of the following operations:

Volume mount

A library client sends a request to the library manager for access to a particular volume in the shared library device. For a scratch volume, the library client does not specify a volume name. If the library manager cannot access the requested

volume, or if scratch volumes are unavailable, the library manager denies the mount request. If the mount is successful, the library manager returns the name of the drive where the volume is mounted.

Volume release

When a library client no longer needs to access a volume, it notifies the library manager that the volume can be returned to a scratch volume. The library manager database is updated with the new location for the volume, which is now in the inventory of the library server. The volume is deleted from the volume inventory of the library client.

Table 1 shows the interaction between library clients and the library manager in processing IBM Spectrum Protect operations.

Table 1. How SAN-enabled servers process IBM Spectrum Protect operations

Operation (Command)	Library manager	Library client
Query library volumes (QUERY LIBVOLUME)	Displays the volumes that are checked into the library. For private volumes, the owner server is also displayed.	Not applicable.
Check in and check out library volumes (CHECKIN LIBVOLUME, CHECKOUT LIBVOLUME)	Sends the commands to the library device.	Not applicable. When a check-in operation is required because of a client restore operation, a request is sent to the library manager server.
Move media and move DRM media (MOVE MEDIA, MOVE DRMEDIA)	Valid only for volumes that are used by the library manager server.	Requests that the library manager server completes the operation. Generates a check-out process on the library manager server.
Audit library inventory (AUDIT LIBRARY)	Synchronizes the inventory with the library device.	Synchronizes the inventory with the library manager server.
Label a library volume (LABEL LIBVOLUME)	Labels and checks in volumes.	Not applicable.
Dismount a volume (DISMOUNT VOLUME)	Sends the request to the library device.	Requests that the library manager server completes the operation.
Query a volume (QUERY VOLUME)	Checks whether the volume is owned by the requesting library client and checks whether the volume is in the library device.	Requests that the library manager server completes the operation.

Managing server requests for volumes

IBM Spectrum Protect™ displays requests and status messages to all administrative command-line clients that are started in console mode. These request messages often have a time limit. Successful server operations must be completed within the time limit that is specified; otherwise, the operation times out.

About this task

For automated libraries, use the CHECKIN LIBVOLUME and LABEL LIBVOLUME commands to insert cartridges into slots. If you specify a value for the WAITTIME parameter, a reply message is displayed. If the value of the parameter is 0, no reply is required. When you issue the CHECKOUT LIBVOLUME command, you must insert cartridges into slots and, in all cases, a reply message is displayed.

Procedure

The following table provides information about how to handle different server media tasks.

Task	Details
Use the administrative client for mount messages	<p>The server sends mount request status messages to the server console and to all administrative command-line clients in mount mode or console mode.</p> <p>To start an administrative command-line client in mount mode, issue the <code>dsmadm -mountmode</code> command on the administrative command-line client.</p>
Receive messages about automated libraries	<p>You can view mount messages and error messages about automated libraries on administrative command-line clients in mount mode or console mode. Mount messages are sent to the library and not to an operator. Messages about problems with the library are sent to the mount message queue.</p>
Get information about pending operator requests	<p>To get information about pending operator requests, issue the <code>QUERY REQUEST</code> command or view the mount message queue on an administrative command-line client that is started in mount mode. When you issue the <code>QUERY REQUEST</code> command, the server displays requested actions and the amount of time that is remaining before the requests time out.</p>
Reply to operator requests	<p>When the server requires an explicit reply to a completed mount request, use the <code>REPLY</code> command.</p> <p>The <code>request_number</code> parameter specifies the request identification number that tells the server which pending operator request is completed. This three-digit number is always displayed as part of the request message.</p>
Cancel an operator request	<p>To cancel a mount request for a library, issue the <code>CANCEL REQUEST</code> command. For most requests that are associated with automated SCSI libraries, an operator must complete a hardware or system action to cancel the requested mount. For such requests, the <code>CANCEL REQUEST</code> command is not accepted by the server.</p> <p>The <code>CANCEL REQUEST</code> command must include the request identification number. This number is included in the request message.</p> <p>If you want to mark the requested volume as <code>UNAVAILABLE</code>, issue the <code>CANCEL REQUEST</code> command and specify the <code>PERMANENT</code> parameter. If you specify the <code>PERMANENT</code> parameter, the server does not try to mount the requested volume again. This is useful if, for example, the volume is at a remote site or is otherwise unavailable.</p>
Respond to a volume check-in request	<p>If the server cannot find a particular volume to mount in an automated library, the server requests that the operator check in the volume.</p> <p>If the requested volume is available, place the volume in the library and check it in. For more information, see Checking volumes into an automated library.</p> <p>If the requested volume is unavailable, update the access mode of the volume by issuing the <code>UPDATE VOLUME</code> command and specifying the <code>ACCESS=UNAVAILABLE</code> parameter. Then, cancel the check-in request by using the <code>CANCEL REQUEST</code> command. Do not cancel the client process that caused the request. Use the <code>QUERY REQUEST</code> command to obtain the ID of the request that you want to cancel.</p> <p>If you do not respond to the check-in request from the server within the mount-wait period that is specified for the device class for the storage pool, the server marks the volume as unavailable.</p>

Task	Details
Determine which volumes are mounted	For a report about all volumes that are currently mounted for use by the server, issue the QUERY MOUNT command. The report shows which volumes are mounted, which drives accessed them, and whether the volumes are in use.
Dismount idle volumes	<p>When a volume is idle, the server keeps it mounted for a time that is specified by the mount retention parameter for the device class. Using a mount retention value can reduce the access time when volumes are used repeatedly.</p> <p>To dismount an idle volume from the drive where it is mounted, issue the DISMOUNT VOLUME command.</p> <p>For information about setting mount retention times, see Controlling the amount of time that a volume remains mounted.</p>

Related information:

➔ QUERY REQUEST (Query one or more pending mount requests)

Managing tape drives

You can query, update, and delete tape drives. You can also clean tape drives and configure tape drive encryption and data validation.

- Updating drives
You can change the attributes of a drive definition to take a drive offline or reconfigure it.
- Data validation during read/write operations to tape
To validate data and identify data that is corrupted, you can use a feature that is called logical block protection. If you use logical block protection, IBM Spectrum Protect inserts a cyclic redundancy check (CRC) value at the end of each logical block of data while it is written to tape.
- Cleaning tape drives
You can use the server to manage tape-drive cleaning. The server can control how tape drives in SCSI libraries are cleaned.
- Tape drive replacement
If you replace a drive in a tape library that is defined to IBM Spectrum Protect, you must delete the drive and path definitions for the old drive and define the new drive and path.

Updating drives

You can change the attributes of a drive definition to take a drive offline or reconfigure it.

About this task

You can change the following attributes of a drive:

- The element address, if the drive is in a SCSI
- The cleaning frequency
- The drive status: online or offline


Restriction: If a drive is in use, you cannot change the element number or the device name. For instructions about taking drives offline, see Taking tape drives offline.

If a volume is mounted in the drive but the volume is idle, it can be explicitly dismounted. For instructions about dismounting idle volumes, see Managing server requests for volumes.

Procedure

- Change the element address of a drive by issuing the UPDATE DRIVE command. For example, in a library that is named AUTO, change the element address of DRIVE3 to 119 by issuing the following command:

```
update drive auto drive3 element=119
```

- Change the device name of a drive by issuing the UPDATE PATH command. For example, to change the device name of a drive that is named DRIVE3, issue the following command: 

```
update path server1 drive3 srctype=server desttype=drive library=scsilib
device=/dev/rmt0
```

Linux

```
update path server1 drive3 srctype=server desttype=drive library=scsilib
device=/dev/IBMtape0
```

Windows

```
update path server1 drive3 srctype=server desttype=drive library=scsilib
device=mt3.0.0.0
```

- Taking tape drives offline
You can take a tape drive offline while it is in use. For example, you might take a drive offline to complete maintenance.

Related reference:

[UPDATE PATH \(Change a path\)](#)

Related information:

[UPDATE DRIVE \(Update a drive\)](#)

Data validation during read/write operations to tape

To validate data and identify data that is corrupted, you can use a feature that is called logical block protection. If you use logical block protection, IBM Spectrum Protect™ inserts a cyclic redundancy check (CRC) value at the end of each logical block of data while it is written to tape.

With logical block protection, you can identify errors that occur when data is written to tape and during data transfer from the tape drive to IBM Spectrum Protect through the storage area network. Drives that support logical block protection validate data during read and write operations. The IBM Spectrum Protect server validates data during read operations.

If validation by the drive fails during write operations, the failure can indicate that data was corrupted during transfer to tape. In this case, the IBM Spectrum Protect server fails the write operation. You must restart the operation to continue. If validation by the drive fails during read operations, the failure can indicate that the tape media is corrupted. If validation by the IBM Spectrum Protect server fails during read operations, the failure can indicate that data was corrupted during transfer from the tape drive, and the server tries the operation again. If validation fails consistently, the IBM Spectrum Protect server issues an error message that indicates hardware or connection problems.

If logical block protection is disabled on a tape drive, or the drive does not support logical block protection, the IBM Spectrum Protect server can read protected data. However, the data is not validated.

Logical block protection is superior to the CRC validation that you can specify when you define or update a storage pool. When you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after the data is written to tape.

Restrictions:

- You cannot use logical block protection for sequential data such as backup sets and database backups.
- CRC checking impacts performance because more processor usage is required on both the client and server to calculate and compare CRC values.
- For a scratch volume, if you specify logical block protection for read/write operations (LBPROTECT=READWRITE), do not change the parameter value at any time after data is written to the volume. Changing the parameter value during the life of the volume on the IBM Spectrum Protect server is not supported.
- Drives that support logical block protection
Logical block protection is available only for 3592, LTO, and ECARTRIDGE device types. Capable 3592 drives include IBM TS1130, TS1140, and later generations. Capable LTO drives include IBM LTO-5 and supported LTO-6 drives. Capable Oracle StorageTek drives include drives with the T10000C and T10000D format.
- Enabling and disabling logical block protection
You can specify logical block protection for read and write operations, or only for write operations. You can also disable logical block protection. By default, logical block protection is disabled because of performance effects that result from cyclic redundancy check (CRC) validation on the server and the tape drive.
- Read/write operations to volumes with logical block protection
Read/write operations to empty or filling volumes depend on whether the volumes have logical block protection. Protected and unprotected data blocks cannot be mixed on the same volume.

- Storage pool management in a tape library
To mix protected and unprotected data in a library, you must create different device classes and different storage pools to separate the data. If a device class is associated with protected data, you can specify logical block protection for read and write operations or for write operations only.

Drives that support logical block protection

Logical block protection is available only for 3592, LTO, and ECARTRIDGE device types. Capable 3592 drives include IBM TS1130, TS1140, and later generations. Capable LTO drives include IBM LTO-5 and supported LTO-6 drives. Capable Oracle StorageTek drives include drives with the T10000C and T10000D format.

The following table shows the media and the formats that you can use with drives that support logical block protection.

Drive	Tape media	Drive formats
IBM TS1130	3592 Generation 2	3592-3 and 3592-3C
IBM TS1140	3592 Generation 2 3592 Generation 3	Generation 2: 3592-3 and 3592-3C Generation 3: 3592-4 and 3592-4C
IBM TS1150	3592 Generation 3 3592 Generation 4	Generation 4: 3592-5 and 3592-5C
IBM LTO-5	LTO-5	Ultrium 5 and Ultrium 5C
IBM LTO-6	LTO-6 LTO-5	Ultrium 6 and Ultrium 6C Ultrium 5 and Ultrium 5C
IBM LTO-7	LTO-7 LTO-6	Ultrium 7 and Ultrium 7C Ultrium 6 and Ultrium 6C
Oracle T10000C	Oracle StorageTek T10000 T2	T10000C and T10000C-C
Oracle T10000D	Oracle StorageTek T10000 T2	T10000D and T10000D-C

Tips:

- To enable logical block protection for a tape volume and then reuse the volume to back up data, you must enable logical block protection for the device class and the drive.
- If you have a 3592, LTO, or Oracle StorageTek drive that is not capable of logical block protection, you can upgrade the drive with firmware that provides logical block protection.

Logical block protection is available for drives that are in SCSI libraries. For the most current information about support for logical block protection, see technote 1568108.

To use logical block protection for write operations, all drives in the library must support logical block protection. If a drive is not capable of logical block protection, volumes that have read/write access are not mounted. However, the server can use the drive to mount volumes that have read-only access. The protected data is read and validated by the IBM Spectrum Protect™ server if logical block protection is enabled for read/write operations.

Enabling and disabling logical block protection

You can specify logical block protection for read and write operations, or only for write operations. You can also disable logical block protection. By default, logical block protection is disabled because of performance effects that result from cyclic redundancy check (CRC) validation on the server and the tape drive.

About this task

Read/write operations to empty or filling volumes depend on whether the volumes have logical block protection. Protected and unprotected data blocks cannot be mixed on the same volume. If you change the setting for logical block protection, the change applies only to empty volumes. Filling and full volumes maintain their status of logical block protection until they are empty and

ready to be refilled. For example, if you disable logical block protection and the server selects a volume that is associated with a device class that has logical block protection, the server continues writing protected data to the volume.

Restriction: Logical block protection is available only for certain device types. For more information, see [Drives that support logical block protection](#).

Procedure

1. To enable logical block protection for the 3592, LTO, and ECARTRIDGE device types, issue the DEFINE DEVCLASS or the UPDATE DEVCLASS command and specify the LBPROTECT parameter. For example, to specify logical block protection during read and write operations for a 3592 device class that is named 3592_lbprotect, issue the following command:

```
define devclass 3592_lbprotect library=3594 lbprotect=readwrite
```

Tips:

- o If you update the value of the LBPROTECT parameter from NO to READWRITE or WRITEONLY and the server selects a filling volume without logical block protection for write operations, the server issues a message each time the volume is mounted. The message indicates that data is written to the volume without logical block protection. To prevent this message from displaying or to have IBM Spectrum Protect™ write data only with logical block protection, update the access of filling volumes without logical block protection to read-only.
 - o To improve performance, do not specify the CRCDATA parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.
 - o When data is validated during read operations by both the drive and by the IBM Spectrum Protect server, it can slow server performance during restore and retrieve operations. To reduce the time that is required for restore and retrieve operations, change the setting of the LBPROTECT parameter from READWRITE to WRITEONLY. After data is restored or retrieved, you can reset the LBPROTECT parameter to READWRITE.
2. To disable logical block protection, issue the DEFINE DEVCLASS or the UPDATE DEVCLASS command and specify the LBPROTECT=NO parameter.

Restriction: If logical block protection is disabled, the server does not write to an empty tape with logical block protection. However, if a filling volume with logical block protection is selected, the server continues to write to the volume with logical block protection. To prevent the server from writing to tapes with logical block protection, change the access of filling volumes with logical block protection to read-only. When data is read, the CRC results are not checked by the drive or server.

If a disaster occurs and the disaster recovery site does not have drives that support logical block protection, you must specify the LBPROTECT=NO parameter. If the tape drives are used for write operations, you must change the volume access for volumes with protected data to read-only to prevent the server from using the volumes.

If the server must enable logical block protection, the server issues an error message that indicates that the drive does not support logical block protection.

What to do next

To determine whether a volume has logical block protection, issue the QUERY VOLUME command and review the value in the `Logical Block Protection` field.

Related reference:

[DEFINE DEVCLASS](#) (Define a device class)

[UPDATE STGPOOL](#) (Update a storage pool)

Related information:

[DEFINE STGPOOL](#) (Define a volume in a storage pool)

[QUERY VOLUME](#) (Query storage pool volumes)

[UPDATE DEVCLASS](#) (Update a device class)

Read/write operations to volumes with logical block protection

Read/write operations to empty or filling volumes depend on whether the volumes have logical block protection. Protected and unprotected data blocks cannot be mixed on the same volume.

If you use the UPDATE DEVCLASS command to change the setting for logical block protection, the change applies only to empty volumes. Filling and full volumes maintain their status of logical block protection until they are empty and ready to be refilled.

For example, suppose that you change the value of the LBPROTECT parameter from READWRITE to NO. If the server selects a volume that is associated with the device class and that has logical block protection, the server continues writing protected data to the volume.

Tips:

- If a drive does not support logical block protection, volumes with logical block protection for write operations cannot be mounted. To prevent the server from mounting the protected volumes for write operations, change the volume access to read-only. Also, disable logical block protection to prevent the server from enabling the feature on the tape drive.
- If a drive does not support logical block protection, and logical block protection is disabled, the server reads data from protected volumes. However, the data is not validated by the server and the tape drive.

Related information:

- [QUERY VOLUME](#) (Query storage pool volumes)
- [UPDATE DEVCLASS](#) (Update a device class)

Storage pool management in a tape library

To mix protected and unprotected data in a library, you must create different device classes and different storage pools to separate the data. If a device class is associated with protected data, you can specify logical block protection for read and write operations or for write operations only.

To define device classes and storage pools for a TS3500 library that has LTO-5 drives, for protected and unprotected data, you can issue a series of commands as shown in the following example:

```
define library 3584 libtype=scsi
define devclass lbprotect library=3584 devicetype=lto lbprotect=readwrite
define devclass normal library=3584 devicetype=lto lbprotect=no
define stgpool lbprotect_pool lbprotect maxscratch=10
define stgpool normal_pool normal maxscratch=10
```

Related reference:

- [DEFINE DEVCLASS](#) (Define a device class)

Related information:

- [DEFINE LIBRARY](#) (Define a library)
- [DEFINE STGPOOL](#) (Define a volume in a storage pool)

Cleaning tape drives

You can use the server to manage tape-drive cleaning. The server can control how tape drives in SCSI libraries are cleaned.

About this task

You must have system privilege or unrestricted storage privilege to clean tape drives. For automated libraries, you can automate cleaning by specifying the frequency of cleaning operations and checking a cleaner cartridge into the library volume inventory. IBM Spectrum Protect™ mounts the cleaner cartridge as specified. There are special considerations if you plan to use server-controlled drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

Tip: If an automated tape library supports library-drive cleaning, ensure that the feature is enabled.

You can prevent premature wear on the read/write heads of drives by using the library cleaning functions that are available from your device manufacturer.

Drives and libraries from manufacturers differ in how they manage cleaner cartridges, and how they report the presence of a cleaner cartridge in a drive. The device driver might not be able to open a drive that contains a cleaner cartridge. Sense codes and error codes that are issued by devices for drive cleaning vary. Library-drive cleaning is usually not known to applications. Therefore, IBM Spectrum Protect might not always detect the cleaner cartridges in drives and might not be able to determine when cleaning begins.

Some devices require a small amount of idle time between mount requests to start drive cleaning. However, IBM Spectrum Protect tries to minimize the idle time for a drive. The result might be to prevent the library drive cleaning from functioning effectively. If this happens, use IBM Spectrum Protect to control drive cleaning. You can set the frequency to match the cleaning recommendations from the manufacturer.

- **Methods for cleaning tape drives**
Over time, the read heads on tapes can get dirty, which can cause read and write operations to fail. To prevent these issues, enable tape cleaning. You can enable tape cleaning from the drive or from IBM Spectrum Protect.
- **Configuring the server for drive cleaning in an automated library**
When you configure server-controlled drive cleaning in an automated library, you can specify how often you want the drives to be cleaned.
- **Resolving errors that are related to drive cleaning**
While moving cartridges within a library, you might place a data cartridge where a cleaner cartridge should be. Review the process that the server completes and the messages that are issued so that you can resolve the issue.

Methods for cleaning tape drives

Over time, the read heads on tapes can get dirty, which can cause read and write operations to fail. To prevent these issues, enable tape cleaning. You can enable tape cleaning from the drive or from IBM Spectrum Protect™.

You can choose to use the library-drive cleaning method or the IBM Spectrum Protect drive-cleaning method, but not both. Some SCSI libraries provide automatic drive cleaning. Select the library-drive cleaning method if it is available. If it is unavailable or causes issues, use IBM Spectrum Protect to control library drive cleaning.

Library drive-cleaning method

The library drive-cleaning method provides several advantages for automated tape libraries that use this function:

- Reduces the burden on the IBM Spectrum Protect administrator to physically manage cartridge cleaning.
- Improves cleaning cartridge usage rates. Most tape libraries track the number of times that drives can be cleaned based on hardware indicators. IBM Spectrum Protect uses a raw count.
- Reduces unnecessary cleaning. Modern tape drives do not have to be cleaned at fixed intervals and can detect and request when cleaning is required.

Manufacturers who provide a library drive-cleaning method recommend its use to prevent premature wear on the read/write heads of the drives. Drives and libraries from different manufacturers differ in how they manage cleaner cartridges and how they report the presence of a cleaner cartridge in a drive. The device driver might not be able to open a drive that contains a cleaner cartridge. Sense codes and error codes that are issued by devices for drive cleaning vary. Library drive cleaning is usually transparent to all applications. However, IBM Spectrum Protect might not always detect cleaner cartridges in drives and might not be able to determine when cleaning begins.

IBM Spectrum Protect drive cleaning method

Some devices require a small amount of idle time between mount requests to start drive cleaning. However, IBM Spectrum Protect tries to minimize the idle time for a drive. The result might be to prevent the library drive cleaning from functioning effectively. If this happens, try using IBM Spectrum Protect to control drive cleaning. Set the frequency to match the cleaning recommendations from the manufacturer.

If IBM Spectrum Protect controls the drive-cleaning process, disable the library drive-cleaning function to prevent problems. If the library drive-cleaning function is enabled, some devices automatically move any cleaner cartridge that is found in the library to slots in the library that are dedicated to cleaner cartridges. You cannot check a cleaner cartridge into the IBM Spectrum Protect library inventory until you disable the library drive-cleaning function.

To enable cleaning from the drive, follow the instructions that are provided by the drive manufacturer. To enable cleaning by using IBM Spectrum Protect, see [Configuring the server for drive cleaning in an automated library](#).

Configuring the server for drive cleaning in an automated library

When you configure server-controlled drive cleaning in an automated library, you can specify how often you want the drives to be cleaned.

Before you begin

Determine how often the drive must be cleaned. This step is required so that you can specify an appropriate value for the CLEANFREQUENCY parameter on the DEFINE DRIVE or UPDATE DRIVE command. For example, to clean a drive after 100 GB of data is processed on the drive, you would specify CLEANFREQUENCY=100.

For guidelines about cleaning frequency, see the drive manufacturer's documentation. If the documentation provides guidelines for cleaning frequency in terms of hours of use, convert the value to a gigabyte value by completing the following steps:

1. Use the bytes-per-second value for the drive to determine a gigabytes-per-hour value.
2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

You can either specify a value for the CLEANFREQUENCY parameter or specify ASNEEDED to clean the drive as needed.

Restrictions:

1. For IBM® 3592 drives, you must specify a numerical value for the CLEANFREQUENCY parameter. By using the cleaning frequency that is listed in the product documentation, you will not overclean the drives.
2. The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. To determine whether a drive supports this function, see the information for your operating system:
 - o [AIX](#) | [Windows](#) Supported devices for AIX and Windows
 - o [Linux](#) Supported devices for Linux

In the technote, click the drive name to view detailed information. If the ASNEEDED value is not supported, specify the number of gigabytes.

Procedure

To configure server-controlled drive cleaning in an automated library, complete the following steps:

Define or update the drives in the library, by using the CLEANFREQUENCY parameter in the DEFINE DRIVE or UPDATE DRIVE command. For example, to clean a drive that is named DRIVE1 after 100 GB of data is processed, issue the following command:

```
update drive autolib1 drive1 cleanfrequency=100
```

Results

After the cleaner cartridge is checked in, the server mounts the cleaner cartridge in a drive when the drive needs cleaning. The server uses that cleaner cartridge for the number of specified cleanings. For more information, see Operations with cleaner cartridges.

What to do next

Check the cleaner cartridge into the library volume inventory by following the instructions in Checking a cleaner cartridge into a library.

- [Checking a cleaner cartridge into a library](#)
To enable automatic tape-drive cleaning, you must check a cleaner cartridge into the volume inventory of the automated library.
- [Operations with cleaner cartridges](#)
To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Related information:

[DEFINE DRIVE](#) (Define a drive to a library)

[UPDATE DRIVE](#) (Update a drive)

Checking a cleaner cartridge into a library

To enable automatic tape-drive cleaning, you must check a cleaner cartridge into the volume inventory of the automated library.

About this task

When you check a cleaner cartridge into a library, ensure that it is correctly identified to the server as a cleaner cartridge. Ensure that a cleaner cartridge is not in a slot that is detected by the search process. Errors and delays of 15 minutes or more might indicate that a cleaner cartridge is improperly placed.

The preferred method is to check in cleaner cartridges individually. If you have to check in both data cartridges and cleaner cartridges, place the data cartridges in the library and check them in first. Then, check the cleaner cartridge in to the library.

Procedure

To check a cleaner cartridge into a library, issue the CHECKIN LIBVOLUME command. For example, to check in a cleaner cartridge that is named AUTOLIB1, issue the following command:

```
checkin libvolume autolib1 cleanv status=cleaner cleanings=10
checklabel=no
```

The server requests that the cartridge is placed in the entry/exit port, or into a specific slot.

Related reference:

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Operations with cleaner cartridges

To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Monitoring the cleaning process

If a cleaner cartridge is checked in to a library, and a drive must be cleaned, the server dismounts the data volume and runs the cleaning operation. If the cleaning operation fails or is canceled, or if no cleaner cartridge is available, you might not be aware that the drive needs cleaning. Monitor cleaning messages for these problems to ensure that drives are cleaned as needed. If necessary, issue the CLEAN DRIVE command to have the server try the cleaning again, or manually load a cleaner cartridge into the drive.

Using multiple cleaner cartridges

The server uses a cleaner cartridge for the number of cleanings that you specify when you check in the cleaner cartridge. If you check in two or more cleaner cartridges, the server uses only one of the cartridges until the designated number of cleanings for that cartridge is reached. Then, the server uses the next cleaner cartridge. If you check in two or more cleaner cartridges and issue two or more CLEAN DRIVE commands concurrently, the server uses multiple cartridges at the same time and decrements the remaining cleanings on each cartridge.

Related reference:

[AUDIT LIBRARY \(Audit volume inventories in an automated library\)](#)

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[CLEAN DRIVE \(Clean a drive\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

Related information:

[QUERY LIBVOLUME \(Query a library volume\)](#)

Resolving errors that are related to drive cleaning

While moving cartridges within a library, you might place a data cartridge where a cleaner cartridge should be. Review the process that the server completes and the messages that are issued so that you can resolve the issue.

When a drive needs cleaning, the server loads what its database shows as a cleaner cartridge into the drive. The drive then moves to a READY state, and IBM Spectrum Protect™ detects that the cartridge is a data cartridge. The server completes the following steps:

1. The server attempts to read the internal tape label of the data cartridge.
2. The server ejects the cartridge from the drive and moves it back to the home slot of the cleaner cartridge within the library. If the eject operation fails, the server marks the drive offline and issues a message that the cartridge is still in the drive.
3. The server checks out the cleaner cartridge to avoid selecting it for another drive cleaning request. The cleaner cartridge remains in the library but no longer appears in the IBM Spectrum Protect library inventory.
4. By using the internal tape label, the server checks the volume name against the current library inventory, storage pool volumes, and the volume history file.
 - o If the volume name is not found in the library inventory, a data cartridge might be checked in as a cleaner cartridge by mistake. When the volume is checked out, you do not have to take further action.
 - o If the volume name is found in the library inventory, the server issues messages that manual intervention and a library audit are required. To resolve the issue, follow the instructions in Auditing the volume inventory in a library.

Tape drive replacement

If you replace a drive in a tape library that is defined to IBM Spectrum Protect™, you must delete the drive and path definitions for the old drive and define the new drive and path.

Replacing drive and path definitions is required even if you are exchanging one drive for another of the same type, with the same logical address, physical address, SCSI ID, and port number. Device alias names can change when you change your drive connections.

If the new drive is an upgrade that supports a new media format, you might be required to define a new logical library, device class, and storage pool. Procedures for setting up a policy for a new drive in a multiple-drive library vary, depending on the types of drives and media in the library.

- **Deleting tape drives**
You can delete tape drives from a library. For example, you can delete a drive that you no longer use, or a drive that you want to replace.
- **Replacing drives with others of the same type**
To add a drive that supports the same media formats as the drive it replaces, you must define a new drive and path.
- **Migrating data to upgraded drives**
If you upgrade all of the tape drives in a library, you can preserve your existing policy definitions to migrate and expire existing data, and you can use the new drives to store data.

Deleting tape drives

You can delete tape drives from a library. For example, you can delete a drive that you no longer use, or a drive that you want to replace.

Procedure

1. Stop the IBM Spectrum Protect™ server and shut down the operating system.
2. Remove the old drive and follow the manufacturer's instructions to install the new drive.
3. Restart the operating system and the IBM Spectrum Protect server.
4. Delete the path from the server to the drive. For example, to delete a path from SERVER1 to LIB1, issue the following command:

```
delete path server1 lib1 srctype=server desttype=drive
```

5. Delete the drive definition. For example, issue the following command to delete a drive that is named DLT1 from a library device that is named LIB1:

```
delete drive lib1 dlt1
```

Related reference:

- [DELETEDRIVE \(Delete a drive from a library\)](#)
- [DELETEPATH \(Delete a path\)](#)

Replacing drives with others of the same type

To add a drive that supports the same media formats as the drive it replaces, you must define a new drive and path.

About this task

If a library includes only one model of drive and you want to replace a drive, you must replace the drive with the same model drive. If a library includes mixed models of drives and you want to replace a drive, you can replace the drive with any model drive that exists in the library.

Procedure

1. Delete the path and drive definitions for the old drive. For example, to delete a drive that is named DRIVE1 from a library that is named LIB1, enter the following command:

```
delete path server2 drive1 srctype=server desttype=drive library=lib1  
delete drive lib1 drive1
```

2. Power off the library, remove the original drive, replace it with the new drive, and power on the library.

3. Refresh the host system to ensure that the system detects the new drive.
4. Define the new drive and path. For example, to define a new drive, DRIVE2, and a path to it from SERVER2, if you are using the IBM Spectrum Protect™ device driver, enter the following commands:

AIX

```
define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=/dev/mt0
```

Linux

```
define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=/dev/tmscsi/mt0
```

Windows

```
define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=mt3.0.0.1
```

Tip: You can use your existing library, device class, and storage pool definitions.

Related reference:

- [DELETE DRIVE \(Delete a drive from a library\)](#)
- [DELETE PATH \(Delete a path\)](#)

Migrating data to upgraded drives

If you upgrade all of the tape drives in a library, you can preserve your existing policy definitions to migrate and expire existing data, and you can use the new drives to store data.

Before you begin

The following scenario assumes that you already have a primary storage pool for a DISK device class that is named POOL1.

Procedure

1. To migrate data to a storage pool that is created for the new drives, specify the NEXTSTGPOOL parameter. For example, to migrate data from an existing storage pool, POOL1, to the new storage pool, POOL2, issue the following command:

```
update stgpool pool1 nextstgpool=pool2
```

2. Update the management-class definitions to store data in the DISK storage pool by using the UPDATE MGMTCLASS command.

Related reference:

- [UPDATE MGMTCLASS \(Update a management class\)](#)
- [UPDATE STGPOOL \(Update a storage pool\)](#)

Related information:

- [DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Securing the IBM Spectrum Protect server

Secure the IBM Spectrum Protect™ server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.

- **Managing administrators**
An administrator who has system authority can complete any task with the IBM Spectrum Protect server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.
- **Changing password requirements**
You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect.

- Securing the server on the system
Protect the system where the IBM Spectrum Protect server runs to prevent unauthorized access.

Managing administrators

An administrator who has system authority can complete any task with the IBM Spectrum Protect™ server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.

Procedure

Complete the following tasks to modify administrator settings.

Task	Procedure
Add an administrator.	<p>To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:</p> <ol style="list-style-type: none"> Register the administrator and specify Pa\$#\$twO as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$twO</pre> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>
Change administrative authority.	<p>Change the authority level for an administrator, ADMIN1.</p> <ul style="list-style-type: none"> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre> Revoke system authority for the administrator by issuing the following command: <pre>revoke authority admin1 classes=system</pre>
Remove administrators.	<p>Remove an administrator, ADMIN1, from accessing the IBM Spectrum Protect server by issuing the following command:</p> <pre>remove admin admin1</pre>
Temporarily prevent access to the server.	<p>Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.</p>

Related concepts:

Planning for administrator roles

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect™.

About this task

By enforcing password authentication and managing password restrictions, you protect your data and your servers from potential security risks.

Procedure

Complete the following tasks to change password requirements for IBM Spectrum Protect servers.

Table 1. Authentication tasks for IBM Spectrum Protect servers

Task	Procedure
Set a limit for invalid password attempts.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details, and then click the Properties tab. Set the number of invalid attempts in the Invalid sign-on attempt limit field. <p>The default value at installation is 0.</p>
Set a minimum length for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of characters in the Minimum password length field.
Set the expiration period for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of days in the Password common expiration field.
Disable password authentication.	<p>By default, the server automatically uses password authentication. With password authentication, all users must enter a password to access the server.</p> <p>You can disable password authentication only for passwords that authenticate with the server (LOCAL). By disabling password authentication, you increase the security risk for the server.</p>
Set a default authentication method.	<p>Issue the SET DEFAULTAUTHENTICATION command. For example, to use the server as the default authentication method, issue the following command:</p> <pre>set defaultauthentication local</pre> <p>To update one client node to authenticate with the server, include AUTHENTICATION=LOCAL in the UPDATE NODE command:</p> <pre>update node authentication=local</pre>

Securing the server on the system

Protect the system where the IBM Spectrum Protect™ server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the directories for the server database and the server instance. Keep the access settings for these directories that you configured during implementation.

- Restricting user access to the server
Authority levels determine what an administrator can do with the IBM Spectrum Protect server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.

Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.

Before you begin

You must have system or operator privilege to stop and start the IBM Spectrum Protect™ server.

- Stopping the server
Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.
- Starting the server for maintenance or reconfiguration tasks
Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Stopping the server

Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.

About this task

When you issue the HALT command to stop the server, the following actions occur:

- All processes and client node sessions are canceled.
- All current transactions are stopped. (The transactions will be rolled back when the server is restarted.)

Procedure

To prepare the system and stop the server, complete the following steps:

1. Prevent new client node sessions from starting by issuing the DISABLE SESSIONS command:

```
disable sessions all
```

2. Determine whether any client node sessions or processes are in progress by completing the following steps:
 - a. On the Overview page of the Operations Center, view the Activity area for the total numbers of processes and sessions that are currently active. If numbers differ significantly from the usual numbers that are displayed during your daily storage-management routine, view other status indicators in the Operations Center to check whether there is a problem.
 - b. View the graph in the Activity area to compare the amount of network traffic over the following periods:
 - The current period, that is, the most recent 24-hour period
 - The previous period, that is, the 24 hours before the current periodIf the graph for the previous period represents the expected amount of traffic, significant differences on the graph for the current period might indicate a problem.
 - c. On the Servers page, select a server for which you want to view processes and sessions, and click Details. If the server is not registered as a hub or spoke server in the Operations Center, obtain information about processes by using administrative commands. Issue the QUERY PROCESS command to query processes and obtain information about sessions by issuing the QUERY SESSION command.
3. Wait until the client node sessions are completed or cancel them. To cancel processes and sessions, complete the following steps:
 - On the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - Click Cancel.
 - If the server is not registered as a hub or spoke server in the Operations Center, cancel sessions by using administrative commands. Issue the CANCEL SESSION command to cancel a session and cancel processes by using the CANCEL PROCESS command.
Tip: If the process that you want to cancel is waiting for a tape volume to be mounted, the mount request is canceled. For example, if you issue an EXPORT, IMPORT, or MOVE DATA command, the command might initiate a process that requires a tape volume to be mounted. However, if a tape volume is being mounted by an automated library, the cancel operation might not take effect until the mount process is complete. Depending on your system environment, this could take several minutes.
4. Stop the server by issuing the HALT command:

```
halt
```

Starting the server for maintenance or reconfiguration tasks

Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSErv utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```




Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode. Follow the instructions for your operating system:
 - o  Starting the server instance
 - o  Starting the server instance
 - o  Starting the server instance

Operations that were disabled during maintenance mode are reenabled.

Planning to upgrade the server

When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect™ server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.

About this task

Follow these guidelines:

- The preferred method is to upgrade the server by using the installation wizard. After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.
- If upgrades are available for both the server component and the Operations Center component, select the check boxes to upgrade both components.

Procedure

1. Review the list of fix packs and interim fixes. See technote 1239415.
2. Review product improvements, which are described in readme files.
Tip: When you obtain the installation package file from the IBM Spectrum Protect support site, you can also access the readme file.
3. Ensure that the version that you upgrade your server to is compatible with other components, such as storage agents and library clients. See technote 1302789.
4. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted. See technote 1053218.
5. Review the upgrade instructions. Ensure that you back up the server database, the device configuration information, and the volume history file.

What to do next

To install a fix pack or interim fix, follow the instructions for your operating system:

- **AIX** Installing an IBM Spectrum Protect server fix pack
- **Linux** Installing an IBM Spectrum Protect server fix pack
- **Windows** Installing an IBM Spectrum Protect server fix pack

Related information:

[Upgrade and Migration Process - Frequently Asked Questions](#)

Preparing for an outage or system update

Prepare IBM Spectrum Protect™ to maintain your system in a consistent state during a planned power outage or system update.

About this task

Ensure that you schedule activities regularly to manage, protect, and maintain the server. For information about scheduling activities such as backing up the database, backing up the device configuration file, and backing up the volume history, see [Defining schedules for server maintenance activities](#).

Procedure

1. Cancel processes and sessions that are in progress by completing the following steps:
 - a. In the Operations Center, on the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - b. Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - c. Click Cancel.
2. Stop the server by issuing the HALT command:

```
halt
```

Tip: You can issue the halt command from the Operations Center by hovering over the Settings icon and clicking Command Builder. Then, select the server, type `halt`, and press Enter.

Related reference:

[HALT \(Shut down the server\)](#)

Preparing for and recovering from a disaster by using DRM

IBM Spectrum Protect™ provides a disaster recovery manager (DRM) function to recover your server and client data during a disaster.

DRM tracks the movement of offsite media and registers that information in the IBM Spectrum Protect database. DRM consolidates plans, scripts, and other information in a plan file that is required to recover the IBM Spectrum Protect server when a disaster or unplanned outage occurs. If you are concerned about possible malware attacks, including ransomware, consider using DRM because it can help you recover your servers after an attack.

Restriction: DRM is only available in the IBM Spectrum Protect Extended Edition product.

- Disaster recovery plan file
The disaster recovery plan file contains the information that is required to recover an IBM Spectrum Protect server to the point in time of the last database backup operation that was completed before the plan was created.
- Recovering the server and client data by using DRM
Use the disaster recovery manager (DRM) function to recover the IBM Spectrum Protect server and client data when a disaster occurs.
- Running a disaster recovery drill
Schedule disaster recovery drills to prepare for audits that certify the recoverability of the IBM Spectrum Protect server and to ensure that data can be restored and operations can resume after an outage. A drill also helps you ensure that all data can be restored and operations resumed before a critical situation occurs.
- Restoring the database
If you have the disaster recovery manager (DRM) function enabled and you followed the procedure to prepare for a disaster, you can restore the database after a disaster. If you do not have DRM configured, you can still restore the database, provided that you have the required backup files.

Disaster recovery plan file

The disaster recovery plan file contains the information that is required to recover an IBM Spectrum Protect™ server to the point in time of the last database backup operation that was completed before the plan was created.

The plan is organized into stanzas, which you can separate into multiple files. Each stanza has a begin statement and an end statement.

Table 1. Stanzas in the disaster recovery plan file

Stanza	Information in the stanza
SERVER.REQUIREMENTS	Identifies the database and recovery log storage requirements for the server.
RECOVERY.INSTRUCTIONS.GENERAL	Identifies site-specific instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.GENERAL. The instructions include the recovery strategy, key contact names, an overview of key applications that are backed up by this server, and other relevant recovery instructions.
RECOVERY.INSTRUCTIONS.OFFSITE	Contains instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.OFFSITE. The instructions describe the name and location of the offsite vault, and how to contact the vault administrator (for example, a name and phone number).
RECOVERY.INSTRUCTIONS.INSTALL	Contains instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.INSTALL. The instructions describe how to rebuild the base server and provide the location of the system image backup copies.
RECOVERY.INSTRUCTIONS.DATABASE	Contains instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.DATABASE. The instructions describe how to prepare for the database recovery. For example, you might enter instructions about how to initialize or load the backup volumes for an automated library. No sample of this stanza is provided.
RECOVERY.INSTRUCTIONS.STGPOOL	Contains instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.STGPOOL. The instructions include the names of your software applications and the copy storage pool names that contain the backups of these applications. No sample of this stanza is provided.

Stanza	Information in the stanza
RECOVERY.VOLUMES.REQUIRED	Provides a list of the database backup and copy storage pool volumes that are required to recover the server. A database backup volume is included if it is part of the most recent database backup series. A copy storage pool volume is included if it is not empty and not marked destroyed.
RECOVERY.DEVICES.REQUIRED	Provides details about the devices that are required to read the backup volumes.
RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE	Contains a script with the commands that are required to recover the server.
RECOVERY.SCRIPT.NORMAL.MODE	Contains a script with the commands that are required to restore the server primary storage pools.
DB.STORAGEPATHS	Identifies the directories for the IBM Spectrum Protect database.
LICENSE.REGISTRATION	Contains a macro to register your server licenses.
COPYSTGPOOL.VOLUMES.AVAILABLE	Contains a macro to mark copy storage pool volumes that were moved offsite and then moved back onsite. You can use the information as a guide and issue the administrative commands. Alternatively, copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
COPYSTGPOOL.VOLUMES.DESTROYED	Contains a macro to mark copy storage pool volumes as unavailable if the volumes were onsite at the time of the disaster. These volumes are considered offsite and have not been destroyed in a disaster. You can use the information as a guide and issue the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
PRIMARY.VOLUMES.DESTROYED	Contains a macro to mark primary storage pool volumes as destroyed if the volumes were onsite at the time of disaster. You can use the information as a guide and run the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
PRIMARY.VOLUMES.REPLACEMENT	Contains a macro to identify replacement primary storage pool volumes. You can use the information as a guide and run the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
STGPOOLS.RESTORE	Contains a macro to restore the primary storage pools. You can use the stanza as a guide and run the administrative commands from a command line. You can also copy, modify, and run it to a file. This macro is started by the RECOVERY.SCRIPT.NORMAL.MODE script.
VOLUME.HISTORY.FILE	Contains a copy of the volume history information when the recovery plan was created. The DSMSERV RESTORE DB utility uses the volume history file to determine what volumes are needed to restore the database. The volume history file is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
DEVICE.CONFIGURATION.FILE	Contains a copy of the server device configuration information when the recovery plan was created. The DSMSERV RESTORE DB utility uses the device configuration file to read the database backup volumes. The device configuration file is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.

Stanza	Information in the stanza
DSMSERV.OPT.FILE	Contains a copy of the server options file. This stanza is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
LICENSE.INFORMATION	Contains a copy of the latest license audit results and the server license terms.
MACHINE.GENERAL.INFORMATION	Provides information for the server machine, such as its location, which is needed to rebuild the server machine. This stanza is included in the plan file if the machine information is saved in the database by using the DEFINE MACHINE command and specifying the ADSMSERVER=YES.
MACHINE.RECOVERY.INSTRUCTIONS	Provides the recovery instructions about the server machine. This stanza is included in the plan file if the machine recovery instructions are saved in the database.
MACHINE.RECOVERY.CHARACTERISTICS	Provides the hardware and software characteristics for the server machine. This stanza is included in the plan file if the machine characteristics are saved in the database.
MACHINE.RECOVERY.MEDIA	Provides information about the media that are required for rebuilding the machine that contains the server. This stanza is included in the plan file if recovery media information is saved in the database and it is associated with the machine that contains the server.

Recovering the server and client data by using DRM

Use the disaster recovery manager (DRM) function to recover the IBM Spectrum Protect™ server and client data when a disaster occurs.

Before you begin

IBM Spectrum Protect is set up to use the Secure Sockets Layer (SSL) protocol for client/server authentication. When you start the server, a digital certificate file, cert.kdb, is created as part of the process. This file includes the server's public key, which allows the client to encrypt data. The digital certificate file cannot be stored in the server database because the Global Security Kit (GSKit) requires a separate file in a certain format.

1. Keep backup copies of the cert.kdb, cert.sth, and cert256.arm files.
2. If both the original certificate files and any copies are lost or corrupted, generate new certificate files.

The master encryption key is stored in a new GSKit-managed key database, dsmkeydb.kdb. If the server has an existing master encryption key, the master encryption key is migrated from the dsmserv.pwd file to the key database, dsmkeydb.kdb. Keep backup copies of the dsmkeydb.kdb and dsmkeydb.sth files. You can configure the BACKUP DB command to back up the master encryption key, or you can manually back up the dsmkeydb.kdb and dsmkeydb.sth files yourself. You cannot recover from a disaster without the master encryption key.

1. Keep backup copies of the dsmkeydb.kdb and dsmkeydb.sth files.

Procedure

1. Get the latest recovery plan.
2. Review the recovery steps that are described in the RECOVERY.INSTRUCTIONS.GENERAL stanza of the plan.
3. Separate the stanzas of the plan file into individual files for general preliminary instructions, IBM Spectrum Protect server recovery scripts, and client recovery instructions.
4. Retrieve all required recovery volumes (as listed in the plan) from the vault.
5. Review the device configuration file to ensure that the hardware configuration at the recovery site is the same as the original site. Any differences must be updated in the device configuration file. The following example configuration changes require updates to the configuration information:
 - o Different device names.
 - o For automated libraries, the requirement of manually placing the database backup volumes in the automated library and updating the configuration information to identify the element within the library. This allows the server to locate

the required database backup volumes.

6. Set up replacement hardware for the IBM Spectrum Protect server, including the operating system and the IBM Spectrum Protect base release installation.
7. Run the IBM Spectrum Protect server recovery scripts from the recovery plan. The RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE and RECOVERY.SCRIPT.NORMAL.MODE stanzas contain executable command files that can be used to drive the recovery of the IBM Spectrum Protect server by calling other command files that were generated in the plan. The RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script recovers the server to the point where clients can begin restores directly from the copy storage pool volumes.
8. Restore the primary storage pools by using the RECOVERY.SCRIPT.NORMAL.MODE script.
9. Start client restore operations in order of highest priority, as defined in your high-level planning.

What to do next

The IBM Spectrum Protect server can now be used for normal server operations. Ensure that all required operations are scheduled. For instructions, see [Defining schedules for server maintenance activities](#) and [Scheduling backup and archive operations](#).

Related reference:

[PREPARE \(Create a recovery plan file\)](#)

Related information:

[Repairing and recovering data in directory-container storage pools](#)

Running a disaster recovery drill

Schedule disaster recovery drills to prepare for audits that certify the recoverability of the IBM Spectrum Protect™ server and to ensure that data can be restored and operations can resume after an outage. A drill also helps you ensure that all data can be restored and operations resumed before a critical situation occurs.

Before you begin

Complete the following tasks:

- Schedule activities regularly to manage, protect, and maintain the server. For more information about scheduling activities, see [Defining schedules for server maintenance activities](#). Ensure that you schedule the following tasks:
 - Backing up the database.
 - Moving media offsite.
 - Backing up the device configuration file, the volume history file, and the dsmserv.opt server options file.
 - **Optional:** Issuing the PREPARE command to create the disaster recovery plan file.

Tip:

When you issue the PREPARE command, the IBM Spectrum Protect disaster recovery manager (DRM) function creates one copy of the disaster recovery plan file.

You can manage offsite disaster recovery without using DRM, however, DRM helps to consolidate plans, scripts, and other information that is required during disaster recovery.

Create multiple copies of the plan for safekeeping. For example, keep copies in print, on a USB flash drive, on disk space that is located offsite, or on a remote server. The disaster recovery plan file is moved offsite daily with the tapes. For more information about DRM, see [Preparing for and recovering from a disaster by using DRM](#).

- Configure the following resources at the disaster recovery site:
 1. A recovery IBM Spectrum Protect server. The server at the disaster recovery site must be at the same level as the server on the production site.
 2. A tape library to store the media that is shipped from the production site. For more information about offsite recovery locations, see [Offsite data storage](#).
 3. Disk storage space for the database, archive log, active logs, and storage pools.
 4. Clients to test restore operations.

About this task

Test the disaster recovery plan and the IBM Spectrum Protect server recoverability often, in an environment that is similar to the production environment.

Procedure

1. Ensure that tapes are available onsite. Issue the QUERY LIBVOLUME command to identify volumes that are checked into an automated library.
2. Back up the database to the onsite tapes by completing the following steps:
 - a. On the Servers page of the Operations Center, select the server whose database you want to back up.
 - b. Click Back Up, and follow the instructions in the Back Up Database window.
3. Copy the following files to the home directory of the server at the recovery site:
 - o Disaster recovery plan file
 - o Volume history file
 - o Device configuration file
 - o Optional: dmserv.opt server options file
4. Move the tape to the offsite recovery location.
5. Restore the server database by using the DSMSEV RESTORE DB utility on the recovery server. For more information about restoring the server database, see Restoring the database.
6. Issue the UPDATE VOLUME command and specify the ACCESS=DESTROYED parameter to indicate that an entire volume must be restored.
7. On the recovery server, restore the storage pool volumes by using the RESTORE STGPOOL command.

What to do next

Ensure that you can access the data in the library by auditing a tape volume in the restored storage pool to verify that the data is consistent. Issue the AUDIT VOLUME command to audit a tape volume. For faster performance, audit restored data only.

Related tasks:

Auditing the volume inventory in a library

Related reference:

[AUDIT VOLUME \(Verify database information for a storage pool volume\)](#)

[DSMSEV RESTORE DB \(Restore the database\)](#)

[RESTORE STGPOOL \(Restore storage pool data\)](#)

Restoring the database

If you have the disaster recovery manager (DRM) function enabled and you followed the procedure to prepare for a disaster, you can restore the database after a disaster. If you do not have DRM configured, you can still restore the database, provided that you have the required backup files.

Before you begin

If the database and recovery log directories are lost, re-create them before you run the DSMSEV RESTORE DB server utility.

About this task

You can restore the database to its most current state or to a specified point in time. To recover the database to the time when the database was lost, recover the database to its latest version.

Restrictions:

- To restore the database to its latest version, you must locate the archive log directory. If you cannot locate the directory, you can restore the database only to a point in time.
- You cannot use the Secure Sockets Layer (SSL) protocol for database restore operations.
- If the release level of the database backup is different from the release level of the server that is being restored, you cannot restore the server database. For example, if you are using a Version 8.1 server and you try to restore a V7.1 database, an error occurs.

Procedure

Use the DSMSEV RESTORE DB server utility to restore the database. Depending on the version of the database that you want to restore, choose one of the following methods:

- Restore a database to its latest version. For example, use the following command:

```
dmserv restore db
```

- Restore a database to a point in time. For example, to restore the database to a backup series that was created on 19 April 2017, use the following command:

```
dsmserv restore db todate=04/19/2017
```

Related reference:

[DSMSERV RESTORE DB \(Restore the database\)](#)

Solutions documentation in PDF files

Prebuilt PDF files for IBM Spectrum Protect™ data protection solutions are available for you to download.

The following prebuilt PDF files are available for IBM Spectrum Protect data protection solutions:

- Introduction to Data Protection Solutions
- Single-Site Disk Solution Guide
- Multisite Disk Solution Guide
- Tape Solution Guide

For more prebuilt PDF files for IBM Spectrum Protect server documentation, see the complete list.

IBM Spectrum Protect servers

IBM Spectrum Protect™ servers store and manage backup, archive, and space-managed data for backup-archive clients and other IBM Spectrum Protect and IBM Spectrum Protect Snapshot components.

- **What's new**
Learn about new features and updates for server components in IBM Spectrum Protect Version 8.1.
- **Installing and upgrading**
You can install or upgrade individual or multiple components in your enterprise network. Solution documentation is available to help you select a best practice solution, based on your business needs and then install, configure, monitor, and operate that solution.
- **Configuring servers**
To complete configuration tasks for the server, review available documentation.
- **Server commands, options, and utilities**
Use commands to administer and configure the server, options to customize the server, and utilities to perform special tasks when the server is not running.
- **Server documentation in PDF files**
Prebuilt PDF files for IBM Spectrum Protect documentation are available for you to download.

What's new

Learn about new features and updates for server components in IBM Spectrum Protect™ Version 8.1.

Tip: To view videos about new features and updates, see the Video library.

To read about new features and updates, follow the links in the table.

Release	New features and updates
V8.1.5	<p>Server</p> <ul style="list-style-type: none"> • Lower the cost of cloud-container storage pools by reclaiming space • Manage your storage environment to help you support General Data Protection Regulation compliance strategies • Generate data deduplication statistics for specified nodes and file spaces • Schedule audit operations to identify corrupted files in a storage pool <p>Operations Center Operations Center updates, including ransomware detection</p>

Release	New features and updates
V8.1.4	<p>Server</p> <ul style="list-style-type: none"> • To enhance security, the default minimum length of passwords is increased • Take advantage of automatic certificate exchange between storage agents, library clients, and library manager servers • Optimize security by using certificates with SHA256 signatures • Specify whether to enforce FIPS 140-2 requirements for encryption • Reduce data fragmentation when you move the contents of storage pool containers <p>Operations Center Operations Center updates</p>
V8.1.3	<p>Server</p> <ul style="list-style-type: none"> • Use cloud tiering for long-term data retention • Install IBM Spectrum Protect on Linux Ubuntu Server LTS • Enhance the security of your storage environment • Take measures to help protect your system against ransomware <p>Operations Center Operations Center updates</p>
V8.1.2	<p>Server</p> <ul style="list-style-type: none"> • Back up data to Microsoft Azure, a cloud-based object storage system • Encrypt client data in a directory-container storage pool • Back up a NAS file server to a directory-container storage pool • Install IBM Spectrum Protect on the Linux on Power Systems™ (little endian) operating system • Protect your storage environment with an improved security protocol • Optimize security with the automatically generated master encryption key • Configure a storage environment by using the Tape Solution Guide • Schedule automatic updates for backup-archive clients • Deprecated and discontinued server options, commands, and parameters <p>Operations Center Operations Center updates</p>
V8.1.1	<p>Server</p> <ul style="list-style-type: none"> • Install IBM Spectrum Protect on the Linux on Power Systems (little endian) operating system • Install IBM Spectrum Protect on the Microsoft Windows Server 2016 operating system • Use the Quantum Scalar i6 library • Review resolved issues <p>Operations Center</p> <ul style="list-style-type: none"> • Review resolved issues

Release	New features and updates
V8.1	<p>Server</p> <ul style="list-style-type: none"> • Meet IBM Spectrum Protect • Secure communications by using the TLS 1.2 protocol • Convert a tape storage pool to a container storage pool • Software upgrade for the server database manager • REGISTER NODE command no longer creates an administrative user ID by default • Optimize user authentication to an Active Directory database • Increased flexibility for protecting and reclaiming tape volumes in container-copy storage pools • Supported operating systems • Monitor the system without using SNMP <p>Operations Center Operations Center updates</p>

- Operations Center updates
New features are available in IBM Spectrum Protect Operations Center Version 8.1.5.
- IBM Spectrum Protect server updates
New features and other changes are available in the IBM Spectrum Protect Version 8.1.5 server.
- Release notes for Version 8.1 server components
Release notes are available for V8.1 components.
- Readme files for Version 8.1 server components
Readme files for Version 8.1 fix packs are published in the IBM Software Support website. Updates might be available for server components, including the server itself, device support, and the Operations Center.

Operations Center updates

New features are available in IBM Spectrum Protect™ Operations Center Version 8.1.5.

The following new features are available:

- Security notifications for potential ransomware attacks. After every client backup session, statistics are analyzed for signs of ransomware infection. If signs are present, a warning message is displayed in the Operations Center. You can use the new Security Notifications page to view details for each security notification. This information helps you determine whether the client is infected with ransomware or if the notification is a false positive.
- Reclamation of space in cloud-container storage pools to help reduce storage costs. Fragmentation occurs in cloud-container storage pools when data is deleted or expired. As a result, a cloud container can have occupied but unused space. You can now specify a threshold for reclaiming that space. When you are selecting a reclamation threshold, you can view the estimated space savings that can be achieved. You can also view the estimated number of requests to move data, and the amount of data to be sent and received. You can use these estimates to decide on a reclamation threshold that is most cost effective based on your cloud provider's storage and data movement charges.

For more information about these enhancements, see the Operations Center help.

Related tasks:

Reclaiming cloud containers

Related reference:

Daily checklist

IBM Spectrum Protect server updates

New features and other changes are available in the IBM Spectrum Protect™ Version 8.1.5 server.

- Lower the cost of cloud-container storage pools by reclaiming space
With IBM Spectrum Protect Version 8.1.5, you can use the new cloud reclamation feature to reclaim space in cloud-container storage pools. You can move data from a larger, fragmented cloud container into a smaller, more fully utilized cloud container. In this way, you can help reduce the cost of using object storage for cloud-container storage pools.
- Manage your storage environment to help you support General Data Protection Regulation compliance strategies
The General Data Protection Regulation (GDPR), effective as of 25 May 2018, is designed to harmonize data privacy

requirements across the European Union (EU). You can use existing features of IBM Spectrum Protect and enhancements that are delivered with IBM Spectrum Protect Version 8.1.5 to help you manage your storage environment to support GDPR compliance strategies.

- Generate data deduplication statistics for specified nodes and file spaces
With IBM Spectrum Protect Version 8.1.5, you can generate data deduplication statistics on a regularly scheduled basis for specified nodes, node groups, and file spaces. By using the DEFINE STGRULE command, you can generate statistics at the same time each day, or at specified intervals.
- Schedule audit operations to identify corrupted files in a storage pool
With IBM Spectrum Protect Version 8.1.5, you can schedule audit operations to identify corrupted files in a storage pool.

Lower the cost of cloud-container storage pools by reclaiming space

With IBM Spectrum Protect™ Version 8.1.5, you can use the new cloud reclamation feature to reclaim space in cloud-container storage pools. You can move data from a larger, fragmented cloud container into a smaller, more fully utilized cloud container. In this way, you can help reduce the cost of using object storage for cloud-container storage pools.

Fragmentation occurs in cloud-container storage pools when data is deleted or expired. To reclaim the unused space in a cloud container, you can schedule a daily cloud reclamation operation by issuing the DEFINE STGRULE server command with the ACTIONTYPE=RECLAIM setting. When the unused space in the cloud container reaches the percentage that you specify, the data is moved into a smaller container. You can schedule an ad hoc reclamation operation by issuing the MOVE CONTAINER command with the default setting DEFRAG=YES.

Alternatively, you can use the Operations Center graphical user interface to schedule cloud reclamation operations and to estimate the space savings that can be achieved.

Related tasks:

Reclaiming cloud containers

Related reference:

DEFINE STGRULE (Define a rule for reclaiming cloud-container storage pools)

MOVE CONTAINER (Move a container)

Operations Center updates

Manage your storage environment to help you support General Data Protection Regulation compliance strategies

The General Data Protection Regulation (GDPR), effective as of 25 May 2018, is designed to harmonize data privacy requirements across the European Union (EU). You can use existing features of IBM Spectrum Protect™ and enhancements that are delivered with IBM Spectrum Protect Version 8.1.5 to help you manage your storage environment to support GDPR compliance strategies.

IBM Spectrum Protect Version 8.1.5 delivers enhancements to allow for an audit trail that you can use to track deletions from the server. These enhancements, together with existing IBM Spectrum Protect functions for deleting data on the server and backup-archive client, can help you comply with the GDPR Article 17, "Right to erasure (right to be forgotten)." Also, existing functions, such as Transport Layer Security (TLS), which enables secure communication of data, can help you comply with Article 32, "Security of processing."

For more information about IBM Spectrum Protect functions that can support GDPR compliance strategies, see technote 22014168.

Generate data deduplication statistics for specified nodes and file spaces

With IBM Spectrum Protect™ Version 8.1.5, you can generate data deduplication statistics on a regularly scheduled basis for specified nodes, node groups, and file spaces. By using the DEFINE STGRULE command, you can generate statistics at the same time each day, or at specified intervals.

To generate statistics on an ad hoc basis, you can run the GENERATE DEDUPSTATS command, followed by the QUERY DEDUPSTATS command. Starting in V8.1.5, you can limit the output of both commands by specifying node lists, node groups, and file spaces. In addition, when you run the QUERY DEDUPSTATS command, you can obtain a summary report of statistics for a specified set of nodes, node groups, and file spaces.

Related tasks:

Defining a storage rule to generate data deduplication statistics

Related reference:

DEFINE STGRULE (Define a rule for generating data deduplication statistics)

GENERATE DEDUPSTATS (Generate data deduplication statistics)

QUERY DEDUPSTATS (Query data deduplication statistics)

Schedule audit operations to identify corrupted files in a storage pool

With IBM Spectrum Protect™ Version 8.1.5, you can schedule audit operations to identify corrupted files in a storage pool.

To schedule the audit operations, use the DEFINE STGRULE command with the ACTIONTYPE=AUDIT setting. If you retain the default value of 7 for the DELAY parameter, the audit operation runs weekly at the same time.

Related tasks:

Auditing a storage pool

Related reference:

DEFINE STGRULE (Define a rule for auditing storage pools)

UPDATE STGRULE (Update a rule for auditing a storage pool)

Release notes for Version 8.1 server components

Release notes are available for V8.1 components.

- Release notes for IBM Spectrum Protect server Version 8.1
The IBM Spectrum Protect server V8.1 is available. Compatibility, installation, and other getting-started issues are addressed.
- Release notes for Operations Center Version 8.1
The Operations Center is a web-based interface that you can use to manage your IBM Spectrum Protect environment. The release notes give you access to the product announcement, known issues, system requirements, installation instructions, and updates.
- Release notes for IBM Spectrum Protect device support Version 8.1
IBM Spectrum Protect device support for V8.1 is available. Compatibility, installation, and other getting-started issues are addressed.

Release notes for IBM Spectrum Protect server Version 8.1

The IBM Spectrum Protect™ server V8.1 is available. Compatibility, installation, and other getting-started issues are addressed.

Contents

- Description
- Announcement
- Compatibility with earlier versions
- System requirements
- Installing and upgrading IBM Spectrum Protect
- Updates, limitations, and known problems

Description

IBM Spectrum Protect provides automated, centrally scheduled, policy-managed backup, archive, and space-management capabilities for file servers, workstations, virtual machines, and applications.

An authorized program analysis report (APAR) is a request for the correction of a defect in a supported release of a program supplied by IBM. For a list of resolved APARs, see APARs fixed in IBM Spectrum Protect server Version 8.1.

Announcement

The announcement for the IBM Spectrum Protect V8.1 family of products includes the following information:

- Detailed product description, including descriptions of new functions
- Product-positioning statement

- International compatibility information

To search for the product announcement, complete the following steps:

1. Go to the product announcement website.
2. In the Search for field, enter the product identifier (PID) for your product. The PID for IBM Spectrum Protect is 5725-W98.
3. In the Information Type field, select Announcement letters, and click Search.
4. From the Search in list, select Product Number.
5. Optional: In the Refine Your Search pane on the left side of the window, select the country where you reside.
6. In the Sort by section, select Newest first.

Compatibility with earlier versions

For compatibility with earlier versions, see the IBM Spectrum Protect Server-Client Compatibility and Upgrade Considerations.

System requirements

For system requirement information, see the IBM Spectrum Protect Supported Operating Systems.

Installing and upgrading IBM Spectrum Protect

For server installation instructions, see the procedure for your operating system:

IBM AIX®

Installing the server

Linux

Installing the server

Microsoft Windows

Installing the server

For upgrade instructions, see Upgrading to V8.1.

Updates, limitations, and known problems

Updates describe new product information or new product features that become available after product release. Updates, limitations, and known issues are documented in the form of technotes in the support knowledge base of the IBM® Support Portal. By searching the knowledge base, you can find workarounds or solutions to known issues.

Updates

REGISTER NODE command no longer creates an administrative user ID by default

Beginning with IBM Spectrum Protect V8.1, the REGISTER NODE command does not automatically create an administrative user ID that matches the node name. This product update can affect the process of registering client nodes, including but not limited to IBM Spectrum Protect backup-archive client nodes. In some cases, you might have to create an administrative user ID by specifying the USERID parameter on the REGISTER NODE command. For information about the types of clients that are affected, see technote 7048963.

To search for the latest updates, see Updates for IBM Spectrum Protect V8.1.

Limitations and known problems

At the time of publication, there were no limitations or known problems.

To search for the latest limitations and known problems, which might include additional items, see Limitations and known problems for IBM Spectrum Protect V8.1.

Release notes for Operations Center Version 8.1

The Operations Center is a web-based interface that you can use to manage your IBM Spectrum Protect™ environment. The release notes give you access to the product announcement, known issues, system requirements, installation instructions, and updates.

Contents

- Description

- Announcement
- Compatibility with the IBM Spectrum Protect server
- System requirements
- Installing or upgrading the Operations Center
- Updates, limitations, and known issues

Description

You can use the Operations Center to take the following actions:

- Identify potential issues with your IBM Spectrum Protect environment
- Monitor key aspects of the storage environment: alerts, clients, servers, policies, storage pools, and storage devices
- Register clients
- Add servers to be monitored
- Back up clients, server databases, and storage pools
- Start storage pool migration and reclamation
- Assign alerts to administrators, and close alerts
- View and cancel server processes and client sessions
- Change client, server, storage pool, and storage device settings
- Create and manage client schedules and view administrative schedules
- Convert primary storage pools to container storage pools
- Copy data from directory-container storage pools to tape
- Configure replication
- Modify policy settings
- Decommission clients and deactivate data
- Create email reports
- View front-end and back-end capacity usage to monitor license compliance
- Issue commands to IBM Spectrum Protect servers

An authorized program analysis report (APAR) is a request for the correction of a defect in a supported release of a program supplied by IBM. For a list of resolved APARs, see APARs fixed in IBM Spectrum Protect Operations Center Version 8.1.

Announcement

The Operations Center is part of the IBM Spectrum Protect V8.1 family of products. The announcement for these products includes the following information:

- Detailed product description, including descriptions of new functions
- Product-positioning statement
- International compatibility information

To search for the product announcement, complete the following steps:

1. Go to the product announcement website.
2. In the Search for field, enter the product identifier (PID) for your product. The PID for IBM Spectrum Protect is 5725-W98.
3. In the Information Type field, select Announcement letters, and click Search.
4. From the Search in list, select Product Number.
5. Optional: In the Refine Your Search pane on the left side of the window, select the country where you reside.
6. In the Sort by section, select Newest first.

Compatibility with the IBM Spectrum Protect server

For compatibility information, see IBM Spectrum Protect server and Operations Center compatibility.

System requirements

For system requirements, see IBM Spectrum Protect Operations Center software and hardware requirements.

Installing or upgrading the Operations Center

For installation instructions, or to upgrade an existing version of the Operations Center, see Installing and upgrading the Operations Center.

Updates, limitations, and known issues

Updates describe new product information or new product features that become available after product release. Updates, limitations, and known issues are documented in the form of technotes in the support knowledge base of the IBM® Support Portal. By searching the knowledge base, you can find workarounds or solutions to known issues.

Updates

For the latest list of updates, see Search results for updates to Operations Center V8.1.

Limitations and known issues

- For a list of limitations and known issues, see Limitations and known issues with Operations Center V8.1.
- To search for additional issues that might become known after the product release, see Search results for known issues with Operations Center V8.1.

Release notes for IBM Spectrum Protect device support Version 8.1

IBM Spectrum Protect™ device support for V8.1 is available. Compatibility, installation, and other getting-started issues are addressed.

Contents

- Description
- Announcement
- Supported devices
- Device driver requirements
- Library information
- Updates, limitations, and known problems

Description

This document includes information about IBM Spectrum Protect V8.1 device drivers.

An authorized program analysis report (APAR) is a request for the correction of a defect in a supported release of a program supplied by IBM. For a list of resolved APARs, see APARs Fixed in IBM Spectrum Protect device driver Version 8.1.

Announcement

IBM Spectrum Protect device support for V8.1 is announced as part of the IBM Spectrum Protect family of products announcement. The announcement for these products includes the following information:

- Detailed product description, including descriptions of new functions
- Product-positioning statement
- International compatibility information

To search for the product announcement, complete the following steps:

1. Go to the product announcement website.
2. In the Search for field, enter the product identifier (PID) for your product. The PID for IBM Spectrum Protect is 5725-W98.
3. In the Information Type field, select Announcement letters, and click Search.
4. From the Search in list, select Product Number.
5. Optional: In the Refine Your Search pane on the left side of the window, select the country where you reside.
6. In the Sort by section, select Newest first.

Supported devices

For information about supported devices and hardware for IBM AIX® and Microsoft Windows systems, see Supported devices for AIX and Windows.

For information about supported devices and hardware for Linux systems, see Supported devices for Linux.

Device driver requirements

Host bus adapter requirements

For best results, connect tape drives and tape libraries to the system on their own host bus adapter. Do not share the host bus adapter with other device types, such as DISK or CD.

Maximum number of devices that are supported by IBM Spectrum Protect device drivers

For information about the maximum number of devices that IBM Spectrum Protect device drivers can support on each operating system, see technote 1364225.

Serial Attached SCSI (SAS) device support

SAS devices can be used on some operating systems and architectures. For information about operating systems and architectures for SAS devices, see technote 1396706.

Running the IBM Spectrum Protect passthru driver with a non-root user ID on Linux operating systems

For information about how a non-root user can use devices with the IBM Spectrum Protect passthru driver on Linux, see technote 1321130. Use option `-g` or `-a` of the device autoconf utility to ensure that non-root users can use devices that are configured with the IBM Spectrum Protect passthru driver. Use option `-g` to add read and write permissions for groups to the SCSI generic driver (`sg`) device files. Use option `-a` to add read and write permissions for all users to the `sg` device files.

Library information

- IBM Spectrum Protect Extended Edition is required for a library with greater than four drives or greater than 48 storage slots.
- The element addresses of the storage slots might not directly correspond to the storage slot numbers. This fact is important because the IBM Spectrum Protect server always references storage slots by element addresses, not storage slot numbers. For element addresses, see the library configuration page for each library.
- For a library with multiple drives, a drive element address is required for the `DEFINE` and `UPDATE DRIVE` commands. However, if the library reports drive serial numbers, you can specify `ELEMENT=AUTODETECT`, and the element address is not required.
- For the procedure to configure the autochanger and each drive in the library separately, see [Configuring and managing storage devices](#).

Updates, limitations, and known problems

Updates

Some devices that were supported by previous releases of IBM Spectrum Protect are no longer supported by the IBM Spectrum Protect V8.1 server. For the latest list of supported devices, see the following links:

- [Supported devices for AIX and Windows](#)
- [Supported devices for Linux](#)

To search for the latest updates, limitations, and known problems, which might include additional items, see [Updates, limitations, and known problems for IBM Spectrum Protect V8.1 device support](#).

Readme files for Version 8.1 server components

Readme files for Version 8.1 fix packs are published in the IBM Software Support website. Updates might be available for server components, including the server itself, device support, and the Operations Center.

View IBM Spectrum Protect™ server V8.1 fix pack readme files

Installing and upgrading

- **Implementing an IBM Spectrum Protect solution**
If you are deploying a new IBM Spectrum Protect server environment, consider implementing a best practice configuration.
- **Installing and upgrading the server**
The IBM Spectrum Protect server provides backup, archive, and space management services to clients. You can install or upgrade individual or multiple servers in your enterprise network.
- **Installing and upgrading the Operations Center**
The Operations Center is the web-based interface for managing your storage environment.

Implementing an IBM Spectrum Protect solution

If you are deploying a new IBM Spectrum Protect™ server environment, consider implementing a best practice configuration.

IBM Spectrum Protect solution documentation is available to help you select a best practice solution, based on your business needs and then install, configure, monitor, and operate that solution.

For details, see [Selecting an IBM Spectrum Protect solution](#).

Availability of features by operating system

Most IBM Spectrum Protect™ features are available on all operating systems that are supported for the server.

In the following table, a check mark indicates that a feature is available.

Table 1. Availability of IBM Spectrum Protect features by operating system

Feature	IBM® AIX®	Linux x86_64	Linux on System z®	Linux on Power Systems™ (little endian)	Microsoft Windows
Aspera® Fast Adaptive Secure Protocol (FASP®) technology: Optimize data transfer to a remote server.		✓ ¹			
Cloud storage by using Amazon Simple Storage Service (Amazon S3) technology.	✓	✓		✓	✓
Cloud storage by using IBM Cloud Object Storage technology.	✓	✓		✓	✓
Cloud storage by using IBM Cloud technology.	✓	✓		✓	✓
Cloud storage by using Microsoft Azure technology.	✓	✓		✓	✓
Cloud storage by using OpenStack Swift technology.	✓	✓		✓	✓
Data deduplication: Use <i>inline data deduplication</i> to eliminate duplicate data while the data is written to a directory-container storage pool or cloud-container storage pool. By using inline data deduplication, you reduce the need for offline reorganization and can improve server performance and lower the cost of storage hardware.	✓	✓	✓	✓	✓
Data deduplication: Use <i>postprocess data deduplication</i> to eliminate duplicate data from sequential-access disk storage pools. This option can result in longer processing times because the server must identify the data, and then remove it from the storage pool.	✓	✓	✓	✓	✓
Disaster recovery manager (DRM): Prepare a plan for recovering your server and client data if a disaster occurs.	✓	✓	✓	✓	✓

Feature	IBM® AIX®	Linux x86_64	Linux on System z®	Linux on Power Systems™ (little endian)	Microsoft Windows
Inline data compression: Compress data as it is written to a cloud-container or directory-container storage pool to reduce the amount of space that the data occupies.	✓	✓	✓	✓	✓
Lightweight Directory Access Protocol (LDAP) authentication: Authenticate users to an Active Directory database on an LDAP server.	✓	✓	✓	✓	✓
Node replication: Incrementally copy data that belongs to backup-archive client nodes from one server to another.	✓	✓	✓	✓	✓
Operations Center: Monitor and manage the storage environment by using the Operations Center, a web-based user interface.	✓	✓	✓	✓	✓
Protection of directory-container storage pools: Protect data in directory-container storage pools by using the PROTECT STGPPOOL command. You can store a copy of the data in another directory-container storage pool on a target replication server, or store a copy on tape in a container-copy storage pool on the same server.	✓	✓ ²	✓	✓	✓
Storage pool encryption: Encrypt data in cloud-container storage pools.	✓	✓		✓	✓
Storage pool encryption: Encrypt data in directory-container storage pools.	✓	✓	✓	✓	✓
Tape storage: Store data on tape, which provides a flexible and affordable option for long-term data retention.	✓	✓ ³	✓	✓	✓
Transport Layer Security (TLS) 1.2 protocol: Secure communications by using TLS 1.2.	✓	✓	✓	✓	✓

¹ Aspera FASP technology is not supported on the Ubuntu Server LTS operating system.

² Protection of directory-container storage pools to tape by using the PROTECT STGPPOOL command is not supported on Ubuntu Server LTS.

³ Tape storage is not supported on Ubuntu Server LTS.

Installing and upgrading the server

The IBM Spectrum Protect™ server provides backup, archive, and space management services to clients. You can install or upgrade individual or multiple servers in your enterprise network.

- Installing the server on AIX systems
- Installing the server on Linux systems

- Installing the server on Windows systems
- Upgrading the server

AIX: Installing the server

Installation of the server includes planning, installation, and initial configuration.

- **AIX: Planning to install the server**
Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.
- **AIX: Installing the server components**
To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.
- **AIX: Taking the first steps after you install IBM Spectrum Protect**
After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect instance.
- **AIX: Installing an IBM Spectrum Protect server fix pack**
IBM Spectrum Protect maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.
- **AIX: Reverting from Version 8.1.5 to a previous server**
If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect server with minimal loss of data.
- **AIX: Reference: DB2 commands for IBM Spectrum Protect server databases**
Use this list as reference when you are directed to issue DB2® commands by IBM® support.
- **AIX: Uninstalling IBM Spectrum Protect**
You can use the following procedures to uninstall IBM Spectrum Protect. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

AIX: Planning to install the server

Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.

- **AIX: What you should know first**
Before installing IBM Spectrum Protect, be familiar with your operating systems, storage devices, communication protocols, and system configurations.
- **AIX: Planning for optimal performance**
Before you install the IBM Spectrum Protect server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.
- **AIX** **AIX: Minimum system requirements for AIX systems**
Before you install an IBM Spectrum Protect server on an AIX operating system, review the hardware and software requirements.
- **AIX** **AIX: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system**
You can install other products that deploy and use DB2® products on the same system as the IBM Spectrum Protect Version 8.1.5 server, with some limitations.
- **AIX: IBM Installation Manager**
IBM Spectrum Protect uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.
- **AIX: Worksheets for planning details for the server**
You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect server. You can also use them to keep track of names and user IDs.
- **AIX: Capacity planning**
Capacity planning for IBM Spectrum Protect includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.
- **AIX: Server naming best practices**
Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect server.

- AIX: Installation directories
Installation directories for the IBM Spectrum Protect server include the server, DB2, device, language, and other directories. Each one contains several additional directories.

AIX: What you should know first

Before installing IBM Spectrum Protect™, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

Server maintenance releases, client software, and publications are available from the IBM® Support Portal.

AIX Restriction: You can install and run the Version 8.1.5 server on a system that already has DB2® installed on it, whether DB2 was installed independently or as part of some other application, with some restrictions. For details, see the compatibility with other DB2 products topic.

Experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. Do not, however, use DB2 tools to change DB2 configuration settings from those that are preset by IBM Spectrum Protect, or alter the DB2 environment for IBM Spectrum Protect in other ways, such as with other products. The V8.1.5 server has been built and tested extensively using the data definition language (DDL) and database configuration that the server deploys.

Attention: Do not alter the DB2 software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

AIX: Planning for optimal performance

Before you install the IBM Spectrum Protect™ server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.

Procedure

1. Review AIX: What you should know first.
2. Review each of the following sub-sections.
 - AIX: Planning for the server hardware and the operating system
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - AIX: Planning for the server database disks
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - AIX: Planning for the server recovery log disks
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - AIX: Planning for directory-container and cloud-container storage pools
Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.
 - AIX: Planning for storage pools in DISK or FILE device classes
Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.
 - AIX: Planning for the correct type of storage technology
Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect.
 - AIX: Applying best practices to the server installation
Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

AIX: Planning for the server hardware and the operating system

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Does the operating system and hardware meet or exceed requirements?</p> <ul style="list-style-type: none"> • Number and speed of processors • System memory • Supported operating system level 	<p>If you are using the minimum required amount of memory, you can support a minimal workload.</p> <p>You can experiment by adding more system memory to determine whether the performance is improved. Then, decide whether you want to keep the system memory dedicated to the server. Test the memory variations by using the entire daily cycle of the server workload.</p> <p>If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.</p> <p>AIX Restriction: Do not use Active Memory™ Expansion (AME). When you use AME, DB2® software uses 4 KB pages instead of 64 KB pages. Each 4 KB page must be decompressed when accessed, and compressed when not needed. When the compression or decompression occurs, DB2 and the server wait for access to the page, which degrades the server performance.</p>	<p>Review operating system requirements at technote 1243309.</p> <p>Additionally, review the guidance in Tuning tasks for operating systems and other applications.</p> <p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For more information about sizing requirements for the server and storage, see the IBM Spectrum Protect™ Blueprint.</p>
<p>Are disks configured for optimal performance?</p>	<p>The amount of tuning that can be done for different disk systems varies. Ensure that the appropriate queue depths and other disk system options are set.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"

Question	Tasks, characteristics, options, or settings	More information
Does the server have enough memory?	<p>Heavier workloads and advanced features such as data deduplication and node replication require more than the minimum system memory that is specified in the system requirements document.</p> <p>For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements:</p> <ul style="list-style-type: none"> • For databases less than 500 GB, you need 16 GB of memory. • For databases with a size of 500 GB - 1 TB, you need 24 GB of memory. • For databases with a size of 1 TB - 1.5 TB, you need 32 GB of memory. • For databases greater than 1.5 TB, you need 40 GB of memory. <p>Ensure that you allocate extra space for the active log and the archive log for replication processing.</p>	<p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication • Memory requirements
Does the system have enough host bus adapters (HBAs) to handle the data operations that the IBM Spectrum Protect server must run simultaneously?	<p>Understand what operations require use of HBAs at the same time.</p> <p>For example, a server must store 1 GB/sec of backup data while also doing storage pool migration that requires 0.5 GB/sec capacity to complete. The HBAs must be able to handle all of the data at the speed required.</p>	See Tuning HBA capacity.
Is network bandwidth greater than the planned maximum throughput for backups?	<p>Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments.</p> <p>For node replication, network bandwidth must be greater than the planned maximum throughput.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Tuning network performance • Checklist for node replication

Question	Tasks, characteristics, options, or settings	More information
Are you using a preferred file system for IBM Spectrum Protect server files?	Use a file system that ensures optimal performance and data availability. The server uses direct I/O with file systems that support the feature. Using direct I/O can improve throughput and reduce processor use. For more information about the preferred file system for your operating system, see IBM Spectrum Protect server-supported file systems.	For more information, see Configuring the operating system for disk performance.
Are you planning to configure enough paging space?	<p>Paging space, or swap space, extends the memory that is available for processing. When the amount of free RAM in the system is low, programs or data that is not in use are moved from memory to paging space. This action releases memory for other activities, such as database operations.</p> <p>AIX Use a minimum of 32 GB of paging space or 50% of your RAM, whichever value is larger.</p>	

AIX: Planning for the server database disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
Is the database on fast, low-latency disks?	<p>Do not use the following drives for the IBM Spectrum Protect™ database:</p> <ul style="list-style-type: none"> • Nearline SAS (NL-SAS) • Serial Advanced Technology Attachment (SATA) • Parallel Advanced Technology Attachment (PATA) <p>Do not use internal disks that are included by default in most server hardware.</p> <p>Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.</p> <p>If you plan to use the data deduplication functions of IBM Spectrum Protect, focus on disk performance in terms of I/O operations per second (IOPS).</p>	For more information, see Checklist for data deduplication.
Is the database stored on disks or LUNs that are separate from disks or LUNs that are used for the active log, archive log, and storage pool volumes?	<p>Separation of the server database from other server components helps reduce contention for the same resources by different operations that must run at the same time.</p> <p>Tip: The database and the archive log can share an array when you use solid-state drive (SSD) technology.</p>	
If you are using RAID, do you know how to select the optimal RAID level for your system? Are you defining all LUNs with the same size and type of RAID?	<p>When a system must do large numbers of writes, RAID 10 outperforms RAID 5. However, RAID 10 requires more disks than RAID 5 for the same amount of usable storage.</p> <p>If your disk system is RAID, define all your LUNs with the same size and type of RAID. For example, do not mix 4+1 RAID 5 with 4+2 RAID 6.</p>	
If an option to set the strip size or segment size is available, are you planning to optimize the size when you configure the disk system?	If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.	The block size that is used for the database varies depending on the table space. Most table spaces use 8 KB blocks, but some use 32 KB blocks.

Question	Tasks, characteristics, options, or settings	More information
<p>Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database?</p> <p>Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.</p>	<p>Heavier workloads and use of some features require more database storage paths than the minimum requirements.</p> <p>Server operations such as data deduplication drive a high number of input/output operations per second (IOPS) for the database. Such operations perform better when the database has more directories.</p> <p>For server databases that are larger than 2 TB or are expected to grow to that size, use eight directories.</p> <p>Consider planned growth of the system when you determine how many storage paths to create. The server uses the higher number of storage paths more effectively if the storage paths are present when the server is first created.</p> <p>Use the <i>DB2_PARALLEL_IO</i> variable to force parallel I/O to occur on table spaces that have one container, or on table spaces that have containers on more than one physical disk. If you do not set the <i>DB2_PARALLEL_IO</i> variable, I/O parallelism is equal to the number of containers that are used by the table space. For example, if a table space spans four containers, the level of I/O parallelism that is used is 4.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For help with forecasting growth when the server deduplicates data, see technote 1596944.</p> <p>For the most recent information about database size, database reorganization, and performance considerations for IBM Spectrum Protect servers, see technote 1683633.</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>
<p>Are all directories for the database the same size?</p>	<p>Directories that are all the same size ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching.</p> <p>This guideline also applies if you must add storage paths after the initial configuration of the server.</p>	
<p>Are you planning to raise the queue depth of the database LUNs on AIX® systems?</p>	<p>The default queue depth is often too low.</p>	<p>See Configuring AIX systems for disk performance.</p>

AIX: Planning for the server recovery log disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Are the active log and archive log stored on disks or LUNs that are separate from what is used for the database and storage pool volumes?</p>	<p>Ensure that the disks where you place the active log are not used for other server or system purposes. Do not place the active log on disks that contain the server database, the archive log, or system files such as page or swap space.</p>	<p>Separation of the server database, active log, and archive log helps to reduce contention for the same resources by different operations that must run at the same time.</p>

Question	Tasks, characteristics, options, or settings	More information
Are the logs on disks that have nonvolatile write cache?	Nonvolatile write cache allows data to be written to the logs as fast as possible. Faster write operations for the logs can improve performance for server operations.	
Are you setting the logs to a size that adequately supports the workload?	<p>If you are not sure about the workload, use the largest size that you can.</p> <p>Active log The maximum size is 512 GB, set with the ACTIVELOGSIZE server option.</p> <p>Ensure that there is at least 8 GB of free space on the active log file system after the fixed size active logs are created.</p> <p>Archive log The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Make the archive log at least as large as the active log.</p>	<ul style="list-style-type: none"> For log sizing details, see the recovery log information in technote 1421060. For information about sizing when you use data deduplication, see Checklist for data deduplication.
Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?	The archive failover log is for emergency use by the server when the archive log becomes full. Slower disks can be used for the archive failover log.	<p>Use the ARCHFAILOVERLOGDIRECTORY server option to specify the location of the archive failover log.</p> <p>Monitor the usage of the directory for the archive failover log. If the archive failover log must be used by the server, the space for the archive log might not be large enough.</p>
If you are mirroring the active log, are you using only one type of mirroring?	<p>You can mirror the log by using one of the following methods. Use only one type of mirroring for the log.</p> <ul style="list-style-type: none"> Use the MIRRORLOGDIRECTORY option that is available for the IBM Spectrum Protect™ server to specify a mirror location. Use software mirroring, such as Logical Volume Manager (LVM) on AIX®. Use mirroring in the disk system hardware. 	<p>If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability.</p> <p>For more information, see Configuring and tuning the recovery log.</p>

AIX: Planning for directory-container and cloud-container storage pools

Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
<p>Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the IBM Spectrum Protect™ database?</p>	<p>Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing.</p> <p>Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication), add 1000 IOPS to this minimum.</p> <p>For example, an IBM Spectrum Protect server that is ingesting 3 TB of data per day would need 6000 IOPS for the database disks:</p> $3000 \text{ IOPS minimum} + 3000 (3 \text{ TB} \times 1000 \text{ IOPS}) = 6000 \text{ IOPS}$	<p>For recommendations about disk selection, see "Planning for server database disks".</p> <p>For more information about IOPS, see the IBM Spectrum Protect Blueprints.</p>
<p>Do you have enough memory for the size of your database?</p>	<p>Use a minimum of 40 GB of system memory for IBM Spectrum Protect servers, with a database size of 100 GB, that are deduplicating data. If the retained capacity of backup data grows, the memory requirement might need to be higher.</p> <p>Monitor memory usage regularly to determine whether more memory is required.</p> <p>Use more system memory to improve caching of database pages. The following memory size guidelines are based on the daily amount of new data that you back up:</p> <ul style="list-style-type: none"> • 128 GB of system memory for daily backups of data, where the database size is 1 - 2 TB • 192 GB of system memory for daily backups of data, where the database size is 2 - 4 TB 	<p>Memory requirements</p>

Question	Tasks, characteristics, options, or settings	More information
<p>Have you properly sized the storage capacity for the database active log and archive log?</p>	<p>Configure the server to have a minimum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072.</p> <p>The suggested starting size for the archive log is 1 TB. The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Ensure that there is at least 10% extra disk space for the file system than the size of the archive log.</p> <p>Use a directory for the database archive logs with an initial free capacity of at least 1 TB. Specify the directory by using the ARCHLOGDIRECTORY server option.</p> <p>Define space for the archive failover log by using the ARCHFAILOVERLOGDIRECTORY server option.</p>	<p>For more information about sizing for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Is compression enabled for the archive log and database backups?</p>	<p>Enable the ARCHLOGCOMPRESS server option to save storage space.</p> <p>This compression option is different from inline compression. Inline compression is enabled by default with IBM Spectrum Protect V7.1.5 and later.</p> <p>Restriction: Do not use this option if the amount of backed up data exceeds 6 TB per day.</p>	<p>For more information about compression for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Are the IBM Spectrum Protect database and logs on separate disk volumes (LUNs)?</p> <p>Is the disk that is used for the database configured according to best practices for a transactional database?</p>	<p>The database must not share disk volumes with IBM Spectrum Protect database logs or storage pools, or with any other application or file system.</p>	<p>For more information about server database and recovery log configuration, see Server database and recovery log configuration and tuning.</p>
<p>Are you using a minimum of eight (2.2 GHz or equivalent) processor cores for each IBM Spectrum Protect server that you plan to use with data deduplication?</p>	<p>If you are planning to use client-side data deduplication, verify that client systems have adequate resources available during a backup operation to complete data deduplication processing. Use a processor that is at least the minimum equivalent of one 2.2 GHz processor core per backup process with client-side data deduplication.</p>	<ul style="list-style-type: none"> • Effective planning and use of deduplication • IBM Spectrum Protect Blueprints

Question	Tasks, characteristics, options, or settings	More information
Did you allocate enough storage space for the database?	<p>For a rough estimate, plan for 100 GB of database storage for every 50 TB of data that is to be protected in deduplicated storage pools. <i>Protected data</i> is the amount of data before data deduplication, including all versions of objects stored.</p> <p>As a best practice, define a new container storage pool exclusively for data deduplication. Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.</p>	
Have you estimated storage pool capacity to configure enough space for the size of your environment?	<p>You can estimate capacity requirements for a deduplicated storage pool by using the following technique:</p> <ol style="list-style-type: none"> 1. Estimate the base size of the source data. 2. Estimate the daily backup size by using an estimated change and growth rate. 3. Determine retention requirements. 4. Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements. 5. Apply the deduplication ratio factor. 6. Apply the compression ratio factor. 7. Round up the estimate to consider transient storage pool usage. 	For an example of using this technique, see Effective planning and use of deduplication.
Have you distributed disk I/O over many disk devices and controllers?	<p>Use arrays that consist of as many disks as possible, which is sometimes referred to as wide striping. Ensure that you use one database directory per distinct array on the subsystem.</p> <p>Set the <i>DB2_PARALLEL_IO</i> registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks.</p> <p>When I/O bandwidth is available and the files are large, for example 1 MB, the process of finding duplicates can occupy the resources of an entire processor. When files are smaller, other bottlenecks can occur.</p> <p>Specify eight or more file systems for the deduplicated storage pool device class so that I/O is distributed across as many LUNs and physical devices as possible.</p>	<p>For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes".</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>

Question	Tasks, characteristics, options, or settings	More information
Have you scheduled daily operations based on your backup strategy?	<p>The best practice sequence of operations is in the following order:</p> <ol style="list-style-type: none"> 1. Client backup 2. Storage pool protection 3. Node replication 4. Database backup 5. Expire inventory 	<ul style="list-style-type: none"> • Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools
Do you have enough storage to manage the DB2® lock list?	<p>If you deduplicate data that includes large files or large numbers of files concurrently, the process can result in insufficient storage space. When the lock list storage is insufficient, backup failures, data management process failures, or server outages can occur.</p> <p>File sizes greater than 500 GB that are processed by data deduplication are most likely to deplete storage space. However, if many backup operations use client-side data deduplication, this problem can also occur with smaller-sized files.</p>	For information about tuning the DB2 LOCKLIST parameter, see Tuning server-side data deduplication.
Is sufficient bandwidth available to transfer data to an IBM Spectrum Protect server?	<p>To transfer data to an IBM Spectrum Protect server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required.</p> <p>Use a V7.1.5 server or higher to use inline compression and use a V7.1.6 or later client to enable enhanced compression processing.</p>	For more information, see the enablededup client option.
Have you determined how many storage pool directories to assign to each storage pool?	<p>Assign directories to a storage pool by using the DEFINE STGPOOLDIRECTORY command.</p> <p>Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).</p>	

Question	Tasks, characteristics, options, or settings	More information
<p>Did you allocate enough disk space in the cloud-container storage pool?</p>	<p>To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space:</p> <ul style="list-style-type: none"> • For serial-attached SCSI (SAS) and spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount, in terabytes, for disk space. • Provide 3 TB for flash-based storage systems with fast network connections to on-premises, high-performance cloud systems. • Provide 5 TB for solid-state drive (SSD) systems with fast network connections to high-performance cloud systems. 	
<p>Did you select the appropriate type of local storage?</p>	<p>Ensure that data transfers from local storage to cloud finish before the next backup cycle starts. Tip: Data is removed from local storage soon after it moves to the cloud. Use the following guidelines:</p> <ul style="list-style-type: none"> • Use flash or SSD for large systems that have high-performing cloud systems. Ensure that you have a dedicated 10 GB wide area network (WAN) link with a high-speed connection to the object storage. For example, use flash or SSD if you have a dedicated 10 GB WAN link plus a high-speed connection to either an IBM® Cloud Object Storage location or to an Amazon Simple Storage Service (Amazon S3) data center. • Use larger capacity 15000 rpm SAS disks for these scenarios: <ul style="list-style-type: none"> ◦ Medium-sized systems ◦ Slower cloud connections, for example, 1 GB ◦ When you use IBM Cloud Object Storage as your service provider across several regions • For SAS or spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount for disk space, in terabytes. 	

AIX: Planning for storage pools in DISK or FILE device classes

Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.

Question	Tasks, characteristics, options, or settings	More information
Can the storage pool LUNs sustain throughput rates for 256 KB sequential reads and writes to adequately handle the workload within the time constraints?	<p>When you are planning for peak loads, consider all the data that you want the server to read or write to the disk storage pools simultaneously. For example, consider the peak flow of data from client backup operations and server data-movement operations such as migration that run at the same time.</p> <p>The IBM Spectrum Protect™ server reads and writes to storage pools predominantly in 256 KB blocks.</p> <p>If the disk system includes the capability, configure the disk system for optimal performance with sequential read/write operations rather than random read/write operations.</p>	For more information, see Analyzing the basic performance of disk systems.
Is the disk configured to use read and write cache?	Use more cache for better performance.	
For storage pools that use FILE device classes, have you determined a good size to use for the storage pool volumes?	Review the information in Optimal number and size of volumes for storage pools that use disk. If you do not have the information to estimate a size for FILE device class volumes, start with volumes that are 50 GB.	Typically, problems arise more frequently when the volumes are too small. Few problems are reported when volumes are larger than needed. When you determine the volume size to use, as a precaution choose a size that might be larger than necessary.
For storage pools that use FILE device classes, are you using preallocated volumes?	<p>Scratch volumes can cause file fragmentation.</p> <p>To ensure that a storage pool does not run out of volumes, set the MAXSCRATCH parameter to a value greater than zero.</p>	<p>Use the DEFINE VOLUME server command to preallocate volumes in the storage pool.</p> <p>Use the DEFINE STGPOOL or UPDATE STGPOOL server command to set the MAXSCRATCH parameter.</p>
For storage pools that use FILE device classes, have you compared the maximum number of client sessions to the number of volumes that are defined?	Always maintain enough usable volumes in the storage pools to allow for the expected peak number of client sessions that run at one time. The volumes might be scratch volumes, empty volumes, or partly filled volumes.	For storage pools that use FILE device classes, only one session or process can write to a volume at the same time.

Question	Tasks, characteristics, options, or settings	More information
<p>For storage pools that use FILE device classes, have you set the MOUNTLIMIT parameter of the device class to a value that is high enough to account for the number of volumes that might be mounted in parallel?</p>	<p>For storage pools that use data deduplication, the MOUNTLIMIT parameter is typically in the range of 500 - 1000.</p> <p>Set the value for MOUNTLIMIT to the maximum number of mount points that are needed for all active sessions. Consider parameters that affect the maximum number of mount points that are needed:</p> <ul style="list-style-type: none"> • The MAXSESSIONS server option, which is the maximum number of IBM Spectrum Protect sessions that can run concurrently. • The MAXNUMMP parameter, which sets the maximum number of mount points that each client node can use. <p>For example, if the maximum number of client node backup sessions is typically 100 and each of the nodes has MAXNUMMP=2, multiply 100 nodes by the 2 mount points for each node to get the value of 200 for the MOUNTLIMIT parameter.</p>	<p>Use the REGISTER NODE or UPDATE NODE server command to set the MAXNUMMP parameter for client nodes.</p>
<p>For storage pools that use DISK device classes, have you determined how many storage pool volumes to put on each file system?</p>	<p>How you configure the storage for a storage pool that uses a DISK device class depends on whether you are using RAID for the disk system.</p> <p>If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system.</p> <p>If you are using RAID 5 with $n + 1$ volumes, configure the storage in one of the following ways:</p> <ul style="list-style-type: none"> • Configure n file systems on the LUN and define one storage pool volume per file system. • Configure one file system and n storage pool volumes for the LUN. 	<p>For an example layout that follows this guideline, see Sample layout of server storage pools.</p>
<p>Did you create your storage pools to distribute I/O across multiple file systems?</p>	<p>Ensure that each file system is on a different LUN on the disk system.</p> <p>Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.</p>	<p>For details, see the following topics:</p> <ul style="list-style-type: none"> • Tuning disk storage for the server • Tuning and configuring storage pools and volumes

AIX: Planning for the correct type of storage technology

Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect™.

Procedure

Review the following table to help you to choose the correct type of storage technology for the storage resources that the server requires.

Table 1. Storage technology types for IBM Spectrum Protect storage requirements

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
Solid-state disk (SSD)	Place the database on SSD in the following circumstances: <ul style="list-style-type: none"> You are using IBM Spectrum Protect data deduplication. You are backing up more than 8 TB of new data daily. 	If you place the IBM Spectrum Protect database on an SSD, as a best practice, place the active log on an SSD. If space is not available, use high-performance disk instead.	Save SSDs for use with the database and active log. The archive log and archive failover logs can be placed on slower storage technology types.	Save SSDs for use with the database and active log. Storage pools can be placed on slower storage technology types.
High-performance disk with the following characteristics: <ul style="list-style-type: none"> 15k rpm disk Fibre Channel or serial-attached SCSI (SAS) interface 	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. Isolate the server database from its logs and storage pools, and from data for other applications.	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use high-performance disks for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use high-performance disks for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.
Medium-performance or high-performance disk with the following characteristics: <ul style="list-style-type: none"> 10k rpm disk Fibre Channel or SAS interface 	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. Isolate the server database from its logs and storage pools, and from data for other applications.	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use medium-performance or high-performance disk for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use medium-performance or high-performance disk for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
SATA, network-attached storage	Do not use this storage for the database. Do not place the database on XIV storage systems.	Do not use this storage for the active log.	Use of this slower storage technology is acceptable because these logs are written once and infrequently read.	Use this slower storage technology in the following circumstances: <ul style="list-style-type: none"> • Data is infrequently written, for example written once. • Data is infrequently read.
Tape and virtual tape				Use for long-term retention or if data is infrequently used.

AIX: Applying best practices to the server installation

Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect™ solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Procedure

- The following best practices are the most important for optimal performance and problem prevention.
- Review the table to determine the best practices that apply to your environment.

Best practice	More information
Use fast disks for the server database. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.	Use fast, low-latency disks for the database. Using SSD is essential if you are using data deduplication and node replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: <ul style="list-style-type: none"> ◦ "Planning for server database disks" ◦ "Planning for the correct type of storage technology"
Ensure that the server system has enough memory.	Review operating system requirements in technote 1243309. Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and node replication can require more than the minimum memory that is specified in the system requirements document. If you plan to run multiple instances, each instance requires the memory that is listed for one server. Multiply the memory for one server by the number of instances that are planned for the system.

Best practice	More information
Separate the server database, the active log, the archive log, and disk storage pools from each other.	<p>Keep all IBM Spectrum Protect storage resources on separate disks. Keep storage pool disks separate from the disks for the server database and logs. Storage pool operations can interfere with database operations when both are on the same disks. Ideally, the server database and logs are also separated from each other. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> ○ "Planning for server database disks" ○ "Planning for server recovery log disks" ○ "Planning for storage pools in DISK or FILE device classes"
Use at least four directories for the server database. For larger servers or servers that use advanced features, use eight directories.	<p>Place each directory on a LUN that is isolated from other LUNs and from other applications.</p> <p>A server is considered to be large if its database is larger than 2 TB or is expected to grow to that size. Use eight directories for such servers.</p> <p>See "Planning for server database disks".</p>
If you are using data deduplication, node replication, or both, follow the guidelines for database configuration and other items.	<p>Configure the server database according to the guidelines, because the database is extremely important to how well the server runs when these features are being used. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> ○ Checklist for data deduplication ○ Checklist for node replication
For storage pools that use FILE type device classes, follow the guidelines for the size of storage pool volumes. Typically, 50 GB volumes are best.	<p>Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size.</p> <p>Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements.</p> <p>Isolate the storage devices that are used by IBM Spectrum Protect from other applications that have high I/O, and ensure that there is enough throughput to that storage.</p> <p>For more details, see Checklist for storage pools on DISK or FILE.</p>
Schedule IBM Spectrum Protect client operations and server maintenance activities to avoid or minimize overlap of operations.	<p>For more details, see the following topics:</p> <ul style="list-style-type: none"> ○ Tuning the schedule for daily operations ○ Checklist for server configuration
Monitor operations constantly.	<p>By monitoring, you can find problems early and more easily identify causes. Keep records of monitoring reports for up to a year to help you identify trends and plan for growth. See Monitoring and maintaining the environment for performance.</p>

AIX: Minimum system requirements for AIX® systems

Before you install an IBM Spectrum Protect™ server on an AIX operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Spectrum Protect server installation

The optimal IBM Spectrum Protect environment is set up with data deduplication by using the IBM Spectrum Protect Blueprints.

For the most current information about IBM Spectrum Protect system requirements, see technote 1243309.

AIX: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

You can install other products that deploy and use DB2® products on the same system as the IBM Spectrum Protect™ Version 8.1.5 server, with some limitations.

To install and use other products that use a DB2 product on the same system as the IBM Spectrum Protect server, ensure that the following criteria are met:

Table 1. Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

Criterion	Instructions
Version level	The other products that use a DB2 product must use DB2 version 9 or later. DB2 products include product encapsulation and segregation support beginning with Version 9. Starting with this version, you can run multiple copies of DB2 products, at different code levels, on the same system. For details, see the information about multiple DB2 copies in the DB2 product information.
User IDs and directories	Ensure that the user IDs, fence user IDs, installation location, other directories, and related information are not shared across DB2 installations. Your specifications must be different from the IDs and locations that you used for the IBM Spectrum Protect server installation and configuration. If you used the dsomicgx wizard to configure the server, these are values that you entered when running the wizard. If you used the manual configuration method, review the procedures that you used if necessary to recall the values that were used for the server.
Resource allocation	<p>Consider the resources and capability of the system compared to the requirements for both the IBM Spectrum Protect server and the other applications that use the DB2 product. To provide sufficient resources for the other DB2 applications, you might have to change the IBM Spectrum Protect server settings so that the server uses less system memory and resources. Similarly, if the workloads for the other DB2 applications compete with the IBM Spectrum Protect server for processor or memory resources, the performance of the server in handling the expected client workload or other server operations might be adversely affected.</p> <p>To segregate resources and provide more capability for the tuning and allocation of processor, memory, and other system resources for multiple applications, consider using logical partition (LPAR), workload partition (WPAR), or other virtual workstation support. For example, run a DB2 application on its own virtualized system.</p>

AIX: IBM Installation Manager

IBM Spectrum Protect™ uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install IBM Spectrum Protect. It must remain installed on the system so that IBM Spectrum Protect can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The IBM Spectrum Protect offering contains all of the media that IBM Installation Manager requires to install IBM Spectrum Protect.

Package

The group of software components that are required to install an offering.

The IBM Spectrum Protect package contains the following components:

- IBM Installation Manager installation program

- IBM Spectrum Protect offering

Package group

A set of packages that share a common parent directory.

The default package group for the IBM Spectrum Protect package is `IBM Installation Manager`.

Repository

A remote or local storage area for data and other application resources.

The IBM Spectrum Protect package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of IBM Spectrum Protect.

AIX: Worksheets for planning details for the server

You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect™ server. You can also use them to keep track of names and user IDs.

Item	Space required	Number of directories	Location of directories
The database			
Active log			
Archive log			
Optional: Log mirror for the active log			
Optional: Secondary archive log (failover location for archive log)			

Item	Names and user IDs	Location
The <i>instance user ID</i> for the server, which is the ID you use to start and run the IBM Spectrum Protect server		
The <i>home directory</i> for the server, which is the directory that contains the instance user ID		
The database instance name		
The <i>instance directory</i> for the server, which is a directory that contains files specifically for this server instance (the server options file and other server-specific files)		
The server name, use a unique name for each server		

AIX: Capacity planning

Capacity planning for IBM Spectrum Protect™ includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.

- **AIX: Estimating space requirements for the database**
To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.
- **AIX: Recovery log space requirements**
In IBM Spectrum Protect, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.
- **AIX: Monitoring space utilization for the database and recovery logs**
To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.
- **AIX: Deleting installation rollback files**
You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

AIX: Estimating space requirements for the database

To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.

About this task

Consider using at least 25 GB for the initial database space. Provision file system space appropriately. A database size of 25 GB is adequate for a test environment or a library-manager-only environment. For a production server supporting client workloads, the database size is expected to be larger. If you use random-access disk (DISK) storage pools, more database and log storage space is needed than for sequential-access storage pools.

The maximum size of the IBM Spectrum Protect™ database is 6 TB.

For information about sizing the database in a production environment that is based on the number of files and on storage pool size, see the following topics.

- **AIX: Estimating database space requirements based on the number of files**
If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.
- **AIX: Estimating database space requirements based on storage pool capacity**
To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.
- **AIX: The database manager and temporary space**
The IBM Spectrum Protect server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

AIX: Estimating database space requirements based on the number of files

If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

About this task

To estimate space requirements that is based on the maximum number of files in server storage, use the following guidelines:

- 600 - 1000 bytes for each stored version of a file, including image backups.
Restriction: The guideline does not include space that is used during data deduplication.
- 100 - 200 bytes for each cached file, copy storage pool file, active-data pool file, and deduplicated file.
- Additional space is required for database optimization to support varying data-access patterns and to support server back-end processing of the data. The amount of extra space is equal to 50% of the estimate for the total number of bytes for file objects.

In the following example for a single client, the calculations are based on the maximum values in the preceding guidelines. The examples do not take into account that you might use file aggregation. In general, when you aggregate small files, it reduces the amount of required database space. File aggregation does not affect space-managed files.

Procedure

1. Calculate the number of file versions. Add each of the following values to obtain the number of file versions:
 - a. Calculate the number of backed-up files. For example, as many as 500,000 client files might be backed up at a time. In this example, storage policies are set to keep up to three copies of backed up files:

$$500,000 \text{ files} * 3 \text{ copies} = 1,500,000 \text{ files}$$

- b. Calculate the number of archive files. For example, as many as 100,000 client files might be archived copies.
- c. Calculate the number of space-managed files. For example, as many as 200,000 client files might be migrated from client workstations.

Using 1000 bytes per file, the total amount of database space that is required for the files that belong to the client is 1.8 GB:

$$(1,500,000 + 100,000 + 200,000) * 1000 = 1.8 \text{ GB}$$

2. Calculate the number of cached files, copy storage-pool files, active-data pool files, and deduplicated files:
 - a. Calculate the number of cached copies. For example, caching is enabled in a 5 GB disk storage pool. The high migration threshold of the pool is 90% and the low migration threshold of the pool is 70%. Thus, 20% of the disk pool, or 1 GB, is occupied by cached files. If the average file size is about 10 KB, approximately 100,000 files are in cache at any one time:

$$100,000 \text{ files} * 200 \text{ bytes} = 19 \text{ MB}$$

- b. Calculate the number of copy storage-pool files. All primary storage pools are backed up to the copy storage pool:

$$(1,500,000 + 100,000 + 200,000) * 200 \text{ bytes} = 343 \text{ MB}$$

- c. Calculate the number of active storage-pool files. All the active client-backup data in primary storage pools is copied to the active-data storage pool. Assume that 500,000 versions of the 1,500,000 backup files in the primary storage pool are active:

$$500,000 * 200 \text{ bytes} = 95 \text{ MB}$$

- d. Calculate the number of deduplicated files. Assume that a deduplicated storage pool contains 50,000 files:

$$50,000 * 200 \text{ bytes} = 10 \text{ MB}$$

Based on the preceding calculations, about 0.5 GB of extra database space is required for the client's cached files, copy storage-pool files, active-data pool files, and deduplicated files.

3. Calculate the amount of extra space that is required for database optimization. To provide optimal data access and management by the server, extra database space is required. The amount of extra database space is equal to 50% of the total space requirements for file objects.

$$(1.8 + 0.5) * 50\% = 1.2 \text{ GB}$$

4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

$$1.8 + 0.5 + 1.2 = 3.5 \text{ GB}$$

5. Calculate the total amount of database space that is required for all clients. If the client that was used in the preceding calculations is typical and you have 500 clients, for example, you can use the following calculation to estimate the total amount of database space that is required for all clients:

$$500 * 3.5 = 1.7 \text{ TB}$$

Results

Tip: In the preceding examples, the results are estimates. The actual size of the database might differ from the estimate because of factors such as the number of directories and the length of the path and file names. Periodically monitor your database and adjust its size as necessary.

What to do next

During normal operations, the IBM Spectrum Protect™ server might require temporary database space. This space is needed for the following reasons:

- To hold the results of sorting or ordering that are not already being kept and optimized in the database directly. The results are temporarily held in the database for processing.
- To give administrative access to the database through one of the following methods:
 - A DB2® open database connectivity (ODBC) client
 - An Oracle Java™ database connectivity (JDBC) client
 - Structured Query Language (SQL) to the server from an administrative-client command line

Consider using an extra 50 GB of temporary space for every 500 GB of space for file objects and optimization. See the guidelines in the following table. In the example that is used in the preceding step, a total of 1.7 TB of database space is required for file objects and optimization for 500 clients. Based on that calculation, 200 GB is required for temporary space. The total amount of required database space is 1.9 TB.

Database size	Minimum temporary-space requirement
< 500 GB	50 GB
≥ 500 GB and < 1 TB	100 GB
≥ 1 TB and < 1.5 TB	150 GB
≥ 1.5 and < 2 TB	200 GB
≥ 2 and < 3 TB	250 - 300 GB
≥ 3 and < 4 TB	350 - 400 GB

AIX: Estimating database space requirements based on storage pool capacity

To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.

AIX: The database manager and temporary space

The IBM Spectrum Protect™ server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

The database manager sorts data in a specific sequence, according to the SQL statement that you issue to request the data. Depending on the workload on the server, and if there is more data than the database manager can manage, the data (that is ordered in sequence) is allocated to temporary disk space. Data is allocated to temporary disk space when there is a large result set. The database manager dynamically manages the memory that is used when data is allocated to temporary disk space.

For example, expiration processing can produce a large result set. If there is not enough system memory on the database to store the result set, some of the data is allocated to temporary disk space. During expiration processing, if a node or file space are selected that are too large to process, the database manager cannot sort the data in memory. The database manager must use temporary space to sort data.

To run database operations, consider adding more database space for the following scenarios:

- The database has a small amount of space and the server operation that requires temporary space uses the remaining free space.
- The file spaces are large, or the file spaces have an assigned policy that creates many file versions.
- The IBM Spectrum Protect server must run with limited memory. The database uses the IBM Spectrum Protect server main memory to run database operations. However, if there is insufficient memory available, the IBM Spectrum Protect server allocates temporary space on disk to the database. For example, if 10G of memory is available and database operations require 12G of memory, the database uses temporary space.
- An `out of database space` error is displayed when you deploy an IBM Spectrum Protect server. Monitor the server activity log for messages that are related to database space.

Important: Do not change the DB2 software that is installed with the IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack, of DB2 software to avoid damage to the database.

AIX: Recovery log space requirements

In IBM Spectrum Protect™, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.

- **AIX: Active and archive log space**
When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.
- **AIX: Active-log mirror space**
The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.
- **AIX: Archive-failover log space**
The archive failover log is used by the server if the archive log directory runs out of space.

AIX: Active and archive log space

When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.

In IBM Spectrum Protect™ servers V7.1 and later, the active log can be a maximum size of 512 GB. The archive log size is limited to the size of the file system that it is installed on.

Use the following general guidelines when you estimate the size of the active log:

- The suggested starting size for the active log is 16 GB.
- Ensure that the active log is at least large enough for the amount of concurrent activity that the server typically handles. As a precaution, try to anticipate the largest amount of work that the server manages at one time. Provision the active log with extra space that can be used if needed. Consider using 20% of extra space.
- Monitor used and available active log space. Adjust the size of the active log as needed, depending upon factors such as client activity and the level of server operations.
- Ensure that the directory that holds the active log is as large as, or larger than, the size of the active log. A directory that is larger than the active log can accommodate failovers, if they occur.
- Ensure that the file system that contains the active log directory has at least 8 GB of free space for temporary log movement requirements.

The suggested starting size for the archive log is 48 GB.

The archive log directory must be large enough to contain the log files that are generated since the previous full backup. For example, if you perform a full backup of the database every day, the archive log directory must be large enough to hold the log files for all the client activity that occurs during 24 hours. To recover space, the server deletes obsolete archive log files after a full backup of the database. If the archive log directory becomes full and a directory for archive failover logs does not exist, log files remain in the active log directory. This condition can cause the active log directory to fill up and stop the server. When the server restarts, some of the existing active-log space is released.

After the server is installed, you can monitor archive log utilization and the space in the archive log directory. If the space in the archive log directory fills up, it can cause the following problems:

- The server is unable to perform full database backups. Investigate and resolve this problem.
- Other applications write to the archive log directory, exhausting the space that is required by the archive log. Do not share archive log space with other applications including other IBM Spectrum Protect servers. Ensure that each server has a separate storage location that is owned and managed by that specific server.
- **AIX: Example: Estimating active and archive log sizes for basic client-store operations**
Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.
- **AIX: Example: Estimating active and archive log sizes for clients that use multiple sessions**
If the client option `RESOURCEUTILIZATION` is set to a value that is greater than the default, the concurrent workload for the server increases.

- AIX: Example: Estimating active and archive log sizes for simultaneous write operations
If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.
- AIX: Example: Estimating active and archive log sizes for basic client store operations and server operations
Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.
- AIX: Example: Estimating active and archive log sizes under conditions of extreme variation
Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.
- AIX: Example: Estimating archive log sizes with full database backups
The IBM Spectrum Protect server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.
- AIX: Example: Estimating active and archive log sizes for data deduplication operations
If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

AIX: Example: Estimating active and archive log sizes for basic client-store operations

Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.

To determine the sizes of the active and archive logs for basic client-store operations, use the following calculation:

number of clients x files stored during each transaction
x log space needed for each file

This calculation is used in the example in the following table.

Table 1. Basic client-store operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053 bytes	The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	19.5 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 bytes = 3.5 GB Increase that amount by the suggested starting size of 16 GB: 3.5 + 16 = 19.5 GB

Item	Example values	Description
Archive log: Suggested size	58.5 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement.</p> $3.5 \times 3 = 10.5 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $10.5 + 48 = 58.5 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

AIX: Example: Estimating active and archive log sizes for clients that use multiple sessions

If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.

To determine the sizes of the active and archive logs when clients use multiple sessions, use the following calculation:

$$\text{number of clients} \times \text{sessions for each client} \times \text{files stored during each transaction} \times \text{log space needed for each file}$$

This calculation is used in the example in the following table.

Table 1. Multiple client sessions

Item	Example values		Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	1000	The number of client nodes that back up, archive, or migrate files every night.
Possible sessions for each client	3	3	The setting of the client option RESOURCEUTILIZATION is larger than the default. Each client session runs a maximum of three sessions in parallel.
Files stored during each transaction	4096	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053	3053	<p>The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>

Item	Example values		Description
Active log: Suggested size	26.5 GB ¹	51 GB ¹	<p>The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 3 sessions for each client x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 10.5 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>10.5 + 16 = 26.5 GB</p> <p>The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.</p> <p>(1000 clients x 3 sessions for each client x 4096 files store during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 35 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>35 + 16 = 51 GB</p>
Archive log: Suggested size	79.5 GB ¹	153 GB ¹	<p>Because of the requirement to be able to store archive logs across three server-database backup cycles, the estimate for the active log is multiplied by 3:</p> <p>10.5 x 3 = 31.5 GB</p> <p>35 x 3 = 105 GB</p> <p>Increase those amounts by the suggested starting size of 48 GB:</p> <p>31.5 + 48 = 79.5 GB</p> <p>105 + 48 = 153 GB</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your active log and adjust its size if necessary.</p>			

AIX: Example: Estimating active and archive log sizes for simultaneous write operations

If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.

The log space that is required for each file increases by about 200 bytes for each copy storage pool that is used for a simultaneous write operation. In the example in the following table, data is stored to two copy storage pools in addition to a primary storage pool. The estimated log size increases by 400 bytes for each file. If you use the suggested value of 3053 bytes of log space for each file, the total number of required bytes is 3453.

This calculation is used in the example in the following table.

Table 1. Simultaneous write operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.

Item	Example values	Description
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3453 bytes	3053 bytes plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3453 bytes for each file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size	60 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

AIX: Example: Estimating active and archive log sizes for basic client store operations and server operations

Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.

For example, migration of files from the random-access (DISK) storage pool to a sequential-access disk (FILE) storage pool uses approximately 110 bytes of log space for each file that is migrated. For example, suppose that you have 300 backup-archive clients and each one of them backs up 100,000 files every night. The files are initially stored on DISK and then migrated to a FILE storage pool. To estimate the amount of active log space that is required for the data migration, use the following calculation. The number of clients in the calculation represents the maximum number of client nodes that back up, archive, or migrate files concurrently at any time.

300 clients x 100,000 files for each client x 110 bytes = 3.1 GB

Add this value to the estimate for the size of the active log that calculated for basic client store operations.

AIX: Example: Estimating active and archive log sizes under conditions of extreme variation

Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

AIX: Example: Estimating archive log sizes with full database backups

The IBM Spectrum Protect™ server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.

For example, if a full database backup occurs once a week, the archive log space must be able to contain the information in the archive log for a full week.

The difference in archive log size for daily and full database backups is shown in the example in the following table.

Table 1. Full database backups

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3453 bytes	3053 bytes for each file plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files per transaction x 3453 bytes per file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size with a full database backup every day	60 GB ¹	Because of the requirement to be able to store archive logs across three backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB

Item	Example values	Description
Archive log: Suggested size with a full database every week	132 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. Multiply the result by the number of days between full database backups:</p> $(4 \text{ GB} \times 3) \times 7 = 84 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $84 + 48 = 132 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested starting size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

AIX: Example: Estimating active and archive log sizes for data deduplication operations

If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

The following factors affect requirements for active and archive log space:

The amount of deduplicated data

The effect of data deduplication on the active log and archive log space depends on the percentage of data that is eligible for deduplication. If the percentage of data that can be deduplicated is relatively high, more log space is required.

The size and number of extents

Approximately 1,500 bytes of active log space are required for each extent that is identified by a duplicate-identification process. For example, if 250,000 extents are identified by a duplicate-identification process, the estimated size of the active log is 358 MB:

$$250,000 \text{ extents identified during each process} \times 1,500 \text{ bytes for each extent} = 358 \text{ MB}$$

Consider the following scenario. Three hundred backup-archive clients back up 100,000 files each night. This activity creates a workload of 30,000,000 files. The average number of extents for each file is two. Therefore, the total number of extents is 60,000,000, and the space requirement for the archive log is 84 GB:

$$60,000,000 \text{ extents} \times 1,500 \text{ bytes for each extent} = 84 \text{ GB}$$

A duplicate-identification process operates on aggregates of files. An aggregate consists of files that are stored in a given transaction, as specified by the TXNGROUPMAX server option. Suppose that the TXNGROUPMAX server option is set to the default of 4096. If the average number of extents for each file is two, the total number of extents in each aggregate is 8192, and the space required for the active log is 12 MB:

$$8192 \text{ extents in each aggregate} \times 1500 \text{ bytes for each extent} = 12 \text{ MB}$$

The timing and number of the duplicate-identification processes

The timing and number of duplicate-identification processes also affects the size of the active log. Using the 12 MB active-log size that was calculated in the preceding example, the concurrent load on the active log is 120 MB if 10 duplicate-identification processes are running in parallel:

$$12 \text{ MB for each process} \times 10 \text{ processes} = 120 \text{ MB}$$

File size

Large files that are processed for duplicate identification can also affect the size of the active log. For example, suppose that a backup-archive client backs up an 80 GB, file-system image. This object can have a high number of duplicate extents if, for example, the files included in the file system image were backed up incrementally. For example, assume that a file

system image has 1.2 million duplicate extents. The 1.2 million extents in this large file represent a single transaction for a duplicate-identification process. The total space in the active log that is required for this single object is 1.7 GB:

$$1,200,000 \text{ extents} \times 1,500 \text{ bytes for each extent} = 1.7 \text{ GB}$$

If other, smaller duplicate-identification processes occur at the same time as the duplicate-identification process for a single large object, the active log might not have enough space. For example, suppose that a storage pool is enabled for deduplication. The storage pool has a mixture of data, including many relatively small files that range from 10 KB to several hundred KB. The storage pool also has few large objects that have a high percentage of duplicate extents.

To take into account not only space requirements but also the timing and duration of concurrent transactions, increase the estimated size of the active log by a factor of two. For example, suppose that your calculations for space requirements are 25 GB (23.3 GB + 1.7 GB for deduplication of a large object). If deduplication processes are running concurrently, the suggested size of the active log is 50 GB. The suggested size of the archive log is 150 GB.

The examples in the following tables show calculations for active and archive logs. The example in the first table uses an average size of 700 KB for extents. The example in the second table uses an average size of 256 KB. As the examples show, the average deduplicate-extent size of 256 KB indicates a larger estimated size for the active log. To minimize or prevent operational problems for the server, use 256 KB to estimate the size of the active log in your production environment.

Table 1. Average duplicate-extent size of 700 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	700 KB	700 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average size for extents.
Extents for a given file	1,198,372 bits	6,135,667 bits	Using the average extent size (700 KB), these calculations represent the total number of extents for a given object. The following calculation was used for an 800 GB object: $(800 \text{ GB} \div 700 \text{ KB}) = 1,198,372 \text{ bits}$ The following calculation was used for a 4 TB object: $(4 \text{ TB} \div 700 \text{ KB}) = 6,135,667 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	1.7 GB	8.6 GB	The estimated active log space that are needed for this transaction.

Item	Example values		Description
Active log: Suggested total size	66 GB ¹	79.8 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of two. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 1.7 \text{ GB}) \times 2 = 50 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $50 + 16 = 66 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 8.6 \text{ GB}) \times 2 = 63.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $63.8 + 16 = 79.8 \text{ GB}$
Archive log: Suggested size	198 GB ¹	239.4 GB ¹	<p>Multiply the estimated size of the active log by a factor of 3.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $50 \text{ GB} \times 3 = 150 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $150 + 48 = 198 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $63.8 \text{ GB} \times 3 = 191.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $191.4 + 48 = 239.4 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Table 2. Average duplicate-extent size of 256 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.

Item	Example values		Description
Average size of extents	256 KB	256 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average extent size.
Extents for a given file	3,276,800 bits	16,777,216 bits	<p>Using the average extent size, these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(800 \text{ GB} \div 256 \text{ KB}) = 3,276,800 \text{ bits}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(4 \text{ TB} \div 256 \text{ KB}) = 16,777,216 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	4.5 GB	23.4 GB	The estimated size of the active log space that is required for this transaction.
Active log: Suggested total size	71.6 GB ¹	109.4 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of 2. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 4.5 \text{ GB}) \times 2 = 55.6 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $55.6 + 16 = 71.6 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 23.4 \text{ GB}) \times 2 = 93.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $93.4 + 16 = 109.4 \text{ GB}$

Item	Example values		Description
Archive log: Suggested size	214.8 GB ¹	328.2 GB ¹	<p>The estimated size of the active log multiplied by a factor of 3.</p> <p>The following calculation was used for an 800 GB object:</p> $55.6 \text{ GB} \times 3 = 166.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $166.8 + 48 = 214.8 \text{ GB}$ <p>The following calculation was used for a 4 TB object:</p> $93.4 \text{ GB} \times 3 = 280.2 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $280.2 + 48 = 328.2 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

AIX: Active-log mirror space

The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.

Creating a log mirror is a suggested option. If you increase the size of the active log, the log mirror size is increased automatically. Mirroring the log can affect performance because of the doubled I/O activity that is required to maintain the mirror. The additional space that the log mirror requires is another factor to consider when deciding whether to create a log mirror.

If the mirror log directory becomes full, the server issues error messages to the activity log and to the db2diag.log. Server activity continues.

AIX: Archive-failover log space

The archive failover log is used by the server if the archive log directory runs out of space.

Specifying an archive failover log directory can prevent problems that occur if the archive log runs out of space. If both the archive log directory and the drive or file system where the archive failover log directory is located become full, the data remains in the active log directory. This condition can cause the active log to fill up, which causes the server to halt.

AIX: Monitoring space utilization for the database and recovery logs

To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

Active log

If the amount of available active log space is too low, the following messages are displayed in the activity log:

```
ANR4531I: IC_AUTOBACKUP_LOG_USED_SINCE_LAST_BACKUP_TRIGGER
```

This message is displayed when the active log space exceeds the maximum specified size. The IBM Spectrum Protect™ server starts a full database backup.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR0297I: IC_BACKUP_NEEDED_LOG_USED_SINCE_LAST_BACKUP

This message is displayed when the active log space exceeds the maximum specified size. You must back up the database manually.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR4529I: IC_AUTOBACKUP_LOG_UTILIZATION_TRIGGER

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. If at least one full database backup has occurred, the IBM Spectrum Protect server starts an incremental database backup. Otherwise, the server starts a full database backup.

ANR0295I: IC_BACKUP_NEEDED_LOG_UTILIZATION

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. You must back up the database manually.

Archive log

If the amount of available archive log space is too low, the following message is displayed in the activity log:

ANR0299I: IC_BACKUP_NEEDED_ARCHLOG_USED

The ratio of used archive-log space to available archive-log space exceeds the log utilization threshold. The IBM Spectrum Protect server starts a full automatic database backup.

Database

If the amount of space available for database activities is too low, the following messages are displayed in the activity log:

ANR2992W: IC_LOG_FILE_SYSTEM_UTILIZATION_WARNING_2

The used database space exceeds the threshold for database space utilization. To increase the space for the database, use the EXTEND DBSPACE command, the EXTEND DBSPACE command, or the DSMSERV FORMAT utility with the DBDIR parameter.

ANR1546W: FILESYSTEM_DBPATH_LESS_1GB

The available space in the directory where the server database files are located is less than 1 GB.

When an IBM Spectrum Protect server is created with the DSMSERV FORMAT utility or with the configuration wizard, a server database and recovery log are also created. In addition, files are created to hold database information used by the database manager. The path specified in this message indicates the location of the database information used by the database manager. If space is unavailable in the path, the server can no longer function.

You must add space to the file system or make space available on the file system or disk.

AIX: Deleting installation rollback files

You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

About this task

To delete the files that are no longer needed, use either the installation graphical wizard or the command line in console mode.

- **AIX: Deleting installation rollback files by using a graphical wizard**
You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.
- **AIX: Deleting installation rollback files by using the command line**
You can delete certain installation files that were saved during the installation process by using the command line.

AIX: Deleting installation rollback files by using a graphical wizard

You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.

Procedure

1. Open IBM Installation Manager.

AIX In the directory where IBM Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command to start IBM Installation Manager:

```
./IBMIM
```

2. Click File > Preferences.
3. Select Files for Rollback.
4. Click Delete Saved Files and click OK.

AIX: Deleting installation rollback files by using the command line

You can delete certain installation files that were saved during the installation process by using the command line.

Procedure

1. In the directory where IBM® Installation Manager is installed, go to the following subdirectory:
 - o **AIX** eclipse/tools

For example:

- o **AIX** /opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command to start an IBM Installation Manager command line:
 - o **AIX** ./imcl -c
 3. Enter **P** to select Preferences.
 4. Enter **B** to select Files for Rollback.
 5. Enter **D** to Delete the Files for Rollback.
 6. Enter **A** to Apply Changes and Return to Preferences Menu.
 7. Enter **C** to leave the Preference Menu.
 8. Enter **X** to Exit Installation Manager.

AIX: Server naming best practices

Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect™ server.

Instance user ID

The instance user ID is used as the basis for other names related to the server instance. The instance user ID is also called the instance owner.

For example: tsminst1

The instance user ID is the user ID that must have ownership or read/write access authority to all directories that you create for the database and the recovery log. The standard way to run the server is under the instance user ID. That user ID must also have read/write access to the directories that are used for any FILE device classes.

AIX

Home directory for the instance user ID

The home directory can be created when creating the instance user ID, by using the option (-m) to create a home directory if it does not exist already. Depending on local settings, the home directory might have the form: /home/instance_user_ID

For example: /home/tsminst1

The home directory is primarily used to contain the profile for the user ID and for security settings.

AIX

Database instance name

The database instance name must be the same as the instance user ID under which you run the server instance.

For example: tsminst1

AIX

Instance directory

The instance directory is a directory that contains files specifically for a server instance (the server options file and other server-specific files). It can have any name that you want. For easier identification, use a name that ties the directory to the instance name.

You can create the instance directory as a subdirectory of the home directory for the instance user ID. For example:
`/home/instance_user_ID/instance_user_ID`

The following example places the instance directory in the home directory for user ID tsminst1: `/home/tsminst1/tsminst1`

You can also create the directory in another location, for example: `/tsmsrvr/tsminst1`

The instance directory stores the following files for the server instance:

- The server options file, `dsmserv.opt`
- The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name
- Volume history file, if the `VOLUMEHISTORY` server option does not specify a fully qualified name
- Volumes for `DEVTYPE=FILE` storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)

Database name

The database name is always `TSMDB1`, for every server instance. This name cannot be changed.

Server name

The server name is an internal name for IBM Spectrum Protect, and is used for operations that involve communication among multiple IBM Spectrum Protect servers. Examples include server-to-server communication and library sharing.

AIX The server name is also used when you add the server to the Operations Center so that it can be managed using that interface. Use a unique name for each server. For easy identification in the Operations Center (or from a `QUERY SERVER` command), use a name that reflects the location or purpose of the server. Do not change the name of an IBM Spectrum Protect server after it is configured as a hub or spoke server.

If you use the wizard, the default name that is suggested is the host name of the system that you are using. You can use a different name that is meaningful in your environment. If you have more than one server on the system and you use the wizard, you can use the default name for only one of the servers. You must enter a unique name for each server.

AIX For example:

- `PAYROLL`
- `SALES`

Directories for database space and recovery log

The directories can be named according to local practices. For easier identification, consider using names that tie the directories to the server instance.

For example, for the archive log:

- **AIX** `/tsminst1_archlog`

AIX: Installation directories

Installation directories for the IBM Spectrum Protect™ server include the server, DB2®, device, language, and other directories. Each one contains several additional directories.

The (/opt/tivoli/tsm/server/bin) is the default directory that contains server code and licensing.

The DB2 product that is installed as part of the installation of the IBM Spectrum Protect server has the directory structure as documented in DB2 information sources. Protect these directories and files as you do the server directories. The default directory is /opt/tivoli/tsm/db2.

You can use US English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, and Russian.

AIX: Installing the server components

To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.

About this task

Using the IBM Spectrum Protect™ installation software, you can install the following components:

- server
Tip: The database (DB2®), the Global Security Kit (GSKit) and IBM® Java™ Runtime Environment (JRE) are automatically installed when you select the server component.
- server languages
- license
- devices
- IBM Spectrum Protect for SAN
- Operations Center

AIX Allow approximately 30 - 45 minutes to install a V8.1.5 server, using this guide.

- AIX: Obtaining the installation package
You can obtain the IBM Spectrum Protect installation package from an IBM download site such as Passport Advantage® or IBM Fix Central.
- AIX: Installing IBM Spectrum Protect by using the installation wizard
You can install the server by using the IBM Installation Manager graphical wizard.
- AIX: Installing IBM Spectrum Protect by using console mode
You can install IBM Spectrum Protect by using the command line in console mode.
- AIX: Installing IBM Spectrum Protect in silent mode
You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.
- AIX: Installing server language packages
Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

AIX: Obtaining the installation package

You can obtain the IBM Spectrum Protect™ installation package from an IBM® download site such as Passport Advantage® or IBM Fix Central.

AIX

Before you begin

If you plan to download the files, set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly:

1. To query the maximum file size value, issue the following command:

```
ulimit -Hf
```

2. If the system user limit for maximum file size is not set to unlimited, change it to unlimited by following the instructions in the documentation for your operating system.

Procedure

1. Download the appropriate package file from one of the following websites.
 - o Download the server package from Passport Advantage or Fix Central.
 - o For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you downloaded the package from an IBM download site, complete the following steps:

AIX

- a. Verify that you have enough space to store the installation files when they are extracted from the product package. See the download document for the space requirements:
 - IBM Spectrum Protect technote 4042944
 - IBM Spectrum Protect Extended Edition technote 4042945
 - IBM Spectrum Protect for Data Retention technote 4042946
- b. Download the package file to the directory of your choice. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.
- c. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

- d. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file, for example:

AIX

```
8.1.x.000-IBM-SPSRV-AIX.bin
```

3. **AIX** Ensure that the following command is enabled so that the IBM Spectrum Protect wizards work properly:
 - o **AIX** `lsuser`By default, the command is enabled.
4. Select one of the following methods of installing IBM Spectrum Protect:
 - o AIX: Installing IBM Spectrum Protect by using the installation wizard
 - o AIX: Installing IBM Spectrum Protect by using console mode
 - o AIX: Installing IBM Spectrum Protect in silent mode
5. After you install IBM Spectrum Protect, and before you customize it for your use, go to the IBM Support Portal. Click Support and downloads and apply any applicable fixes.

AIX: Installing IBM Spectrum Protect by using the installation wizard

You can install the server by using the IBM® Installation Manager graphical wizard.

Before you begin

Take the following actions before you start the installation:

- **AIX** If the following RPM files are not installed on your system, you must install them. For instructions, see Installing RPM files for the graphical wizard.
 - o `atk-1.12.3-2.aix5.2.ppc.rpm`
 - o `cairo-1.8.8-1.aix5.2.ppc.rpm`
 - o `expat-2.0.1-1.aix5.2.ppc.rpm`
 - o `fontconfig-2.4.2-1.aix5.2.ppc.rpm`
 - o `freetype2-2.3.9-1.aix5.2.ppc.rpm`
 - o `gettext-0.10.40-6.aix5.1.ppc.rpm`
 - o `glib2-2.12.4-2.aix5.2.ppc.rpm`
 - o `gtk2-2.10.6-4.aix5.2.ppc.rpm`
 - o `libjpeg-6b-6.aix5.1.ppc.rpm`

- o libpng-1.2.32-2.aix5.2.ppc.rpm
- o libtiff-3.8.2-1.aix5.2.ppc.rpm
- o pango-1.14.5-4.aix5.2.ppc.rpm
- o pixman-0.12.0-3.aix5.2.ppc.rpm
- o xcursor-1.1.7-3.aix5.2.ppc.rpm
- o xft-2.1.6-5.aix5.1.ppc.rpm
- o xrender-0.9.1-3.aix5.2.ppc.rpm
- o zlib-1.2.3-3.aix5.1.ppc.rpm
- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install IBM Spectrum Protect™ by using this method:

Option	Description
Installing the software from a downloaded package:	a. Change to the directory where you downloaded the package. b. Start the installation wizard by issuing the following command: <div style="background-color: #800040; color: white; padding: 2px; display: inline-block; margin: 5px 0;">AIX</div> <code>./install.sh</code>

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

You can view installation log files by clicking File > View Log from the Installation Manager tool. To collect these log files, click Help > Export Data for Problem Analysis from the Installation Manager tool.
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- AIX

 After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.
- AIX

AIX: Installing prerequisite RPM files for the graphical wizard
Before you can use the graphical wizard of IBM Installation Manager to install IBM Spectrum Protect, you must ensure that the necessary RPM files are installed.

AIX: Installing IBM Spectrum Protect by using console mode

You can install IBM Spectrum Protect™ by using the command line in console mode.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install IBM Spectrum Protect by using this method:

Option	Description
--------	-------------

Option	Description
Installing the software from a downloaded package:	<p>a. Change to the directory where you downloaded the package.</p> <p>b. Start the installation wizard in console mode by issuing the following command: AIX</p> <pre>./install.sh -c</pre> <p>Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the Summary panel, specify G to generate the responses.</p>

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - **AIX** /var/ibm/InstallationManager/logs
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **AIX** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

AIX: Installing IBM Spectrum Protect in silent mode

You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the input directory where the installation package is extracted:

install_response_sample.xml

Use this file to install the IBM Spectrum Protect™ components.

update_response_sample.xml

Use this file to upgrade the IBM Spectrum Protect components.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
2. If you install the server and Operations Center in silent mode, create a password for the Operations Center truststore in the response file.

If you are using the install_response_sample.xml file, add the password in the following line of the file, where *mypassword* represents the password:

```
<variable name='ssl.password' value='mypassword' />
```

For more information about this password, see Installation checklist

Tip: To upgrade the Operations Center, the truststore password is not required if you are using the update_response_sample.xml file.

3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:

◦ **AIX**

```
./install.sh -s -input response_file -acceptLicense
```

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - **AIX** /var/ibm/InstallationManager/logs
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **AIX** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

AIX

AIX: Installing server language packages

Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Before you begin

For instructions on installing storage agent language packages, see Language pack configuration for storage agents.

- **AIX: Server language locales**
Use either the default language package option or select another language package to display server messages and help.
- **AIX: Configuring a language package**
After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.
- **AIX: Updating a language package**
You can modify or update a language package by using the IBM® Installation Manager.

AIX: Server language locales

Use either the default language package option or select another language package to display server messages and help.

AIX This language package is automatically installed for the following default language option for IBM Spectrum Protect™ server messages and help:

- **AIX** LANGUAGE en_US

For languages or locales other than the default, install the language package that your installation requires.

You can use the languages that are shown:

AIX

Table 1. Server languages for AIX®

Language	LANGUAGE option value
Chinese, Simplified	zh_CN
Chinese, Simplified (UTF-8)	ZH_CN
Chinese, Traditional (Big5)	Zh_TW
Chinese, Traditional (UTF-8)	ZH_TW
Chinese, Traditional (euc_tw)	zh_TW
English	en_US
English (UTF-8)	EN_US
French	fr_FR
French (UTF-8)	FR_FR
German	de_DE
German (UTF-8)	DE_DE
Italian	it_IT
Italian (UTF-8)	IT_IT

Language	LANGUAGE option value
Japanese, EUC	ja_JP
Japanese, PC	Ja_JP
Japanese, UTF8	JA_JP
Korean	ko_KR
Korean (UTF-8)	KO_KR
Portuguese, Brazilian	pt_BR
Portuguese, Brazilian (UTF-8)	PT_BR
Russian	ru_RU
Russian (UTF-8)	RU_RU
Spanish	es_ES
Spanish (UTF-8)	ES_ES

AIX Restriction: For Operations Center users, some characters might not be displayed properly if the web browser does not use the same language as the server. If this problem occurs, set the browser to use the same language as the server.

AIX: Configuring a language package

After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.

About this task

AIX To set support for a certain locale, complete one of the following tasks:

- Set the LANGUAGE option in the server options file to the name of the locale that you want to use. For example:
 - **AIX** To use the `it_IT` locale, set the LANGUAGE option to `it_IT`. See AIX: Server language locales.
- **AIX** If you are starting the server in the foreground, set the `LC_ALL` environment variable to match the value that is set in the server options file. For example, to set the environment variable for Italian, enter the following value:

```
export LC_ALL=it_IT
```

If the locale is successfully initialized, it formats the date, time, and number for the server. If the locale is not successfully initialized, the server uses the US English message files and the date, time, and number format.

AIX: Updating a language package

You can modify or update a language package by using the IBM® Installation Manager.

About this task

You can install another language package within the same IBM Spectrum Protect™ instance.

- Use the Modify function of IBM Installation Manager to install another language package.
- Use the Update function of IBM Installation Manager to update to newer versions of the language packages.

Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.

AIX: Taking the first steps after you install IBM Spectrum Protect

After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect™ instance.

About this task

1. Create the directories and user ID for the server instance. See [AIX: Creating the user ID and directories for the server instance](#).
2. Configure a server instance. Select one of the following options:
 - o Use the configuration wizard, the preferred method. See [AIX: Configuring IBM Spectrum Protect by using the configuration wizard](#).
 - o Manually configure the new instance. See [AIX: Configuring the server instance manually](#). Complete the following steps during a manual configuration.
 - a. Set up your directories and create the IBM Spectrum Protect instance. See [AIX: Creating the server instance](#).
 - b. Create a new server options file by copying the sample file to set up communications between the server and clients. See [AIX: Configuring server and client communications](#).
 - c. Issue the DSMSERV FORMAT command to format the database. See [AIX: Formatting the database and log](#).
 - d. Configure your system for database backup. See [AIX: Preparing the database manager for database backup](#).
3. Configure options to control when database reorganization runs. See [AIX: Configuring server options for server database maintenance](#).
4. Start the server instance if it is not already started.
 - o [AIX](#) See [AIX: Starting the server instance](#).
5. Register your license. See [AIX: Registering licenses](#).
6. Prepare your system for database backups. See [AIX: Preparing the server for database backup operations](#).
7. Monitor the server. See [AIX: Monitoring the server](#).

- [AIX: Creating the user ID and directories for the server instance](#)
Create the user ID for the IBM Spectrum Protect server instance and create the directories that the server instance needs for database and recovery logs.
- [AIX: Configuring the IBM Spectrum Protect server](#)
After you have installed the server and prepared for the configuration, configure the server instance.
- [AIX: Configuring server options for server database maintenance](#)
To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.
- [AIX](#) [AIX: Starting the server instance](#)
You can start the server by using the instance user ID, which is the preferred method, or the root user ID.
- [AIX: Stopping the server](#)
You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.
- [AIX: Registering licenses](#)
Immediately register any IBM Spectrum Protect licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.
- [AIX: Preparing the server for database backup operations](#)
To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.
- [AIX: Running multiple server instances on a single system](#)
You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.
- [AIX: Monitoring the server](#)
When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

AIX: Creating the user ID and directories for the server instance

Create the user ID for the IBM Spectrum Protect™ server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See [AIX: Worksheets for planning details for the server](#).

Procedure

1. Create the user ID that will own the server instance. You use this user ID when you create the server instance in a later step.

AIX

AIX

Create a user ID and group that will be the owner of the server instance.

- a. The following commands can be run from an administrative user ID that will set up the user and group. Create the user ID and group in the home directory of the user.

Restriction: In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID and group name must comply with the following rules:

- The length must be 8 characters or less.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

For example, create user ID `tsminst1` in group `tsmsrvrs`. The following examples show how to create this user ID and group using operating system commands.

AIX

```
mkgroup id=1001 tsmsrvrs
mkuser id=1002 pgrp=tsmsrvrs home=/home/tsminst1 tsminst1
passwd tsminst1
```

Restriction: DB2® does not support direct operating system user authentication through LDAP.

- b. Log off, then log in to your system. Change to the user account that you just created. Use an interactive login program, such as `telnet`, so that you are prompted for the password and can change it if necessary.

2. Create directories that the server requires.

AIX

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	<code>mkdir /tsminst1</code>	
The database directories	<code>mkdir /tsmdb001</code> <code>mkdir /tsmdb002</code> <code>mkdir /tsmdb003</code> <code>mkdir /tsmdb004</code>	
Active log directory	<code>mkdir /tsmlog</code>	
Archive log directory	<code>mkdir /tsmarchlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir /tsmlogmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir /tsmarchlogfailover</code>	

When a server is initially created by using the `DSMSERV FORMAT` utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

AIX: Configuring the IBM Spectrum Protect server

After you have installed the server and prepared for the configuration, configure the server instance.

About this task

Configure an IBM Spectrum Protect™ server instance by selecting one of the following options:

- **AIX: Configuring IBM Spectrum Protect by using the configuration wizard**
The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some

configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect server program.

- AIX: Configuring the server instance manually

After installing IBM Spectrum Protect Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

AIX: Configuring IBM Spectrum Protect by using the configuration wizard

The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect™ server program.

Before you begin

Before you use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing IBM Spectrum Protect, creating the database and log directories, and creating the directories and user ID for the server instance.

Procedure

1. Ensure that the following requirements are met:

AIX

- The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/` directory. Also, ensure that the SSH daemon service has access rights for connecting to the system by using the `localhost` value.
- You must be able to log in to the system with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- A backup copy of the following files must be saved to a safe and secure location:
 - Master encryption key files (`dsmkeydb.*`)
 - Server certificate and private key files (`cert.*`)

2. Start the local version of the wizard:

AIX

- Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be run only by using the root user ID.

Follow the instructions to complete the configuration. The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

AIX: Configuring the server instance manually

After installing IBM Spectrum Protect™ Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

- AIX: Creating the server instance

Create an IBM Spectrum Protect instance by issuing the `db2icrt` command.

- **AIX** AIX: Configuring server and client communications

A default sample server options file, `dsmserv.opt.smp`, is created during IBM Spectrum Protect installation in the `/opt/tivoli/tsm/server/bin` directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.

- AIX: Formatting the database and log

Use the `DSMSERV FORMAT` utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

- AIX: Preparing the database manager for database backup

To back up the data in the database to IBM Spectrum Protect, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

AIX: Creating the server instance

Create an IBM Spectrum Protect™ instance by issuing the db2icrt command.

About this task

You can have one or more server instances on one workstation.

AIX Important: Before you run the db2icrt command, verify the following items:

- The home directory for the user (/home/tsminst1) exists. If there is no home directory, you must create it. The instance directory stores the following files that are generated by the IBM Spectrum Protect server:
 - The server options file, dsmserv.opt
 - The server key database file, cert.kdb, and the .arm files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
 - Device configuration file, if the DEVCONFIG server option does not specify a fully qualified name
 - Volume history file, if the VOLUMEHISTORY server option does not specify a fully qualified name
 - Volumes for DEVTYPE=FILE storage pools, if the directory for the device class is not fully specified, or not fully qualified
 - User exits
 - Trace output (if not fully qualified)
- A backup copy of the following files must be saved to a safe and secure location:
 - Master encryption key files (dsmkeydb.*)
 - Server certificate and private key files (cert.*)
- The root user and instance-user ID must have write permission to the shell configuration file. A shell configuration file (for example, .profile) exists in the home directory. For more information, see the DB2® product information. Search for Linux and UNIX environment variable settings.

AIX

1. Log in using the root user ID and create an IBM Spectrum Protect instance. The name of the instance must be the same name as the user that owns the instance. Use the db2icrt command and enter the command on one line: **AIX**

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
instance_name instance_name
```

For example, if your user ID for this instance is tsminst1, use the following command to create the instance. Enter the command on one line. **AIX**

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
tsminst1 tsminst1
```

Remember: From this point on, use this new user ID when you configure your IBM Spectrum Protect server. Log out of the root user ID and log in under the new instance-user ID.

2. Change the default directory for the database to be the same as the instance directory for the server. If you have multiple servers, log in under the instance ID for each server. Issue this command:

```
db2 update dbm cfg using dftdbpath instance_directory
```

For example, where instance_directory is the instance user ID:

```
db2 update dbm cfg using dftdbpath /tsminst1
```

3. Modify the library path to include libraries that are required for server operations.

Tip: In the following examples, here are the directories:

- *server_bin_directory* is a subdirectory of the server installation directory. For example, /opt/tivoli/tsm/server/bin.
- *instance_users_home_directory* is the home directory of the instance user. For example, /home/tsminst1.

- **AIX** Issue the following command, on one line:

```
export LIBPATH=server_bin_directory/dbbkapi:  
/usr/opt/ibm/gsk8_64/lib64:$LIBPATH
```

- You must update one of the following files to set the library path when DB2 or the server are started. Update per the shell that the instance user is configured to use.

Bash or Korn shell:

```
instance_users_home_directory/sqllib/userprofile
```

C shell:


```
instance_users_home_directory/sqlllib/usercshrc
```

- o Update per the shell that the instance user is configured to use.

Bash or Korn shell:

Add the following entry to the *instance_users_home_directory/sqlllib/userprofile* file, on one line: **AIX**

```
export LIBPATH=server_bin_directory/  
dbbkapi:/usr/opt/ibm/gsk8_64/lib64:$LIBPATH
```

C shell:

Add the following entry to the *instance_users_home_directory/sqlllib/usercshrc* file, on one line: **AIX**

```
setenv LIBPATH server_bin_directory/dbbkapi:  
/usr/opt/ibm/gsk8_64/lib64:$LIBPATH
```

Remember: The following entries must be in the library path, preceding any other entries in the library path:

- server_bin_directory/dbbkapi
- /usr/local/ibm/gsk8_64/lib64

4. Create a new server options file. See AIX: Configuring server and client communications.

AIX

AIX: Configuring server and client communications

A default sample server options file, *dsmserv.opt.smp*, is created during IBM Spectrum Protect™ installation in the */opt/tivoli/tsm/server/bin* directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.

About this task

Ensure that you have a server instance directory, for example */tsminst1*, and copy the sample file to this directory. Name the new file *dsmserv.opt* and edit the options. Complete this set-up before you initialize the server database. Each sample or default entry in the sample options file is a comment, a line beginning with an asterisk (*). Options are not case-sensitive and one or more blank spaces are allowed between keywords and values.

When editing the options file, follow these guidelines:

- Remove the asterisk at the beginning of the line to activate an option.
- Begin entering the options in any column.
- Enter only one option per line, and the option must be on only one line.
- If you make multiple entries for a keyword, the IBM Spectrum Protect server uses the last entry.

If you change the server options file, you must restart the server for the changes to take effect.

You can specify one or more of the following communication methods:

- TCP/IP Version 4 or Version 6
- Shared memory
- Secure Sockets Layer (SSL)

Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the IBM Spectrum Protect server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

- **AIX** AIX: Setting TCP/IP options
Select from a range of TCP/IP options for the IBM Spectrum Protect server or retain the default.
- **AIX** AIX: Setting shared memory options
You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.
- **AIX** AIX: Setting Secure Sockets Layer options
You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

AIX: Setting TCP/IP options

Select from a range of TCP/IP options for the IBM Spectrum Protect™ server or retain the default.

About this task

The following is an example of a list of TCP/IP options that you can use to set up your system.

```
commethod      tcpip
tcpport        1500
tcpwindowsize  0
tcpnodelay     yes
```

Tip: You can use TCP/IP Version 4, Version 6, or both.

TCPPOINT

The server port address for TCP/IP and SSL communication. The default value is 1500.

AIX TCPWINDOWSIZE

AIX Specifies the size of the TCP/IP buffer that is used when sending or receiving data. The window size that is used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but can improve performance.

You can specify an integer from 0 to 2048. To use the default window size for the operating system, specify 0.

TCPNODELAY

Specifies whether or not the server sends small messages or lets TCP/IP buffer the messages. Sending small messages can improve throughput but increases the number of packets sent over the network. Specify YES to send small messages or NO to let TCP/IP buffer them. The default is YES.

TCPADMINPORT

Specifies the port number on which the server TCP/IP communication driver is to wait for TCP/IP or SSL-enabled communication requests other than client sessions. The default is the value of TCPPOINT.

SSLTCPPOINT

(SSL-only) Specifies the Secure Sockets Layer (SSL) port number on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line backup-archive client and the command-line administrative client.

SSLTCPADMINPORT

(SSL-only) Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client.

AIX: Setting shared memory options

You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.

About this task

The following example shows a shared memory setting:

```
commethod      sharedmem
shmport        1510
```

In this example, SHMPORT specifies the TCP/IP port address of a server when using shared memory. Use the SHMPORT option to specify a different TCP/IP port. The default port address is 1510.

COMMETHOD can be used multiple times in the IBM Spectrum Protect™ server options file, with a different value each time. For example, the following example is possible:

```
commethod tcpip
commethod sharedmem
```

AIX The maximum number of concurrent shared memory sessions is based on available system resources. Each shared memory session uses one shared memory region of up to 4 MB, and four IPCS message queues, depending on the IBM Spectrum Protect client level.

AIX If the server and client are not run under the same user ID, then the server must be root. This prevents shared memory communication errors.

AIX: Setting Secure Sockets Layer options

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Before you begin

SSL is the standard technology for creating encrypted sessions between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communication paths. With SSL, the identity of the server is verified through the use of digital certificates.

To ensure better system performance, use SSL only for sessions when it is needed. Consider adding additional processor resources on the IBM Spectrum Protect™ server to manage the increased requirements.

AIX: Formatting the database and log

Use the DSMSEV FORMAT utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

After you set up server communications, you are ready to initialize the database. Ensure that you log in by using the instance user ID. Do not place the directories on file systems that might run out of space. If certain directories (for example, the archive log) become unavailable or full, the server stops.

Setting the exit list handler

Set the DB2NOEXITLIST registry variable to ON for each server instance. Log on to the system as the server instance owner and issue this command:

```
db2set -i server_instance_name DB2NOEXITLIST=ON
```

For example: **AIX**

```
db2set -i tsminst1 DB2NOEXITLIST=ON
```

Initializing a server instance

Use the DSMSEV FORMAT utility to initialize a server instance. For example, if the server instance directory is */tsminst1*, issue the following commands: **AIX**

```
cd /tsminst1
dsmserv format dbdir=/tsmdb001 activelogsiz=32768
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

AIX Tip: If DB2® does not start after you issue the DSMSEV FORMAT command, you might need to disable the file system mount option NOSUID. If this option is set on the file system that contains the DB2 instance owner directory, or on any file system that contains the DB2 database, active logs, archive logs, failover logs, or mirrored logs, the option must be disabled to start the system.

After you disable the NOSUID option, remount the file system and then start DB2 by issuing the following command:

```
db2start
```

Related information:

[DSMSEV FORMAT \(Format the database and log\)](#)

AIX: Preparing the database manager for database backup

To back up the data in the database to IBM Spectrum Protect™, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

About this task

AIX Starting with IBM Spectrum Protect V7.1, it is no longer necessary to set the API password during a manual configuration of the server. If you set the API password during the manual configuration process, attempts to back up the database might fail.

If you use the configuration wizard to create an IBM Spectrum Protect server instance, you do not have to complete these steps. If you are configuring an instance manually, complete the following steps before you issue either the BACKUP DB or the RESTORE DB commands.

Attention: If the database is unusable, the entire IBM Spectrum Protect server is unavailable. If a database is lost and cannot be recovered, it might be difficult or impossible to recover data that is managed by that server. Therefore, it is critically important to back up the database.

AIX In the following commands, replace the example values with your actual values. The examples use `tsminst1` for the server instance user ID, `/tsminst1` for the server instance directory, and `/home/tsminst1` as the server instance users home directory.

1. Set the IBM Spectrum Protect API environment-variable configuration for the database instance:

- a. Log in by using the `tsminst1` user ID.
- b. When user `tsminst1` is logged in, ensure that the DB2® environment is properly initialized. The DB2 environment is initialized by running the `/home/tsminst1/sqllib/db2profile` script, which normally runs automatically from the profile of the user ID. Ensure the `.profile` file exists in the instance users home directory, for example, `/home/tsminst1/.profile`. If `.profile` does not run the `db2profile` script, add the following lines:

```
if [ -f /home/tsminst1/sqllib/db2profile ]; then
    . /home/tsminst1/sqllib/db2profile
fi
```

- c. In the `instance_directory/sqllib/userprofile` file, add the following lines:

```
DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
DSMI_DIR=server_bin_directory/dbbkapi
DSMI_LOG=server_instance_directory
export DSMI_CONFIG DSMI_DIR DSMI_LOG
```

where:

- *instance_directory* is the home directory of the server instance user.
- *server_instance_directory* is the server instance directory.
- *server_bin_directory* is the server bin directory. The default location is `/opt/tivoli/tsm/server/bin`.

In the `instance_directory/sqllib/usercshrc` file, add the following lines:

```
setenv DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
setenv DSMI_DIR=server_bin_directory/dbbkapi
setenv DSMI_LOG=server_instance_directory
```

2. Log off and log in again as `tsminst1`, or issue this command:

```
. ~/.profile
```

Tip: Ensure that you enter a space after the initial dot (.) character.

3. Create a file that is named `tsmdbmgr.opt` in the *server_instance* directory, which is in the `/tsminst1` directory in this example, and add the following line:

```
SERVERNAME TSMDBMGR_TSMINST1
```

Remember: The value for `SERVERNAME` must be consistent in the `tsmdbmgr.opt` and `dsm.sys` files.

4. As root user, add the following lines to the IBM Spectrum Protect API `dsm.sys` configuration file. By default, the `dsm.sys` configuration file is in the following default location:

- o `server_bin_directory/dbbkapi/dsm.sys`

```
servername TSMDBMGR_TSMINST1
commethod tcpip
tcpserveraddr localhost
tcpport 1500
errorlogname /tsminst1/tsmdbmgr.log
nodename $$TSMDBMGR_$$
```

where

- o *servername* matches the *servername* value in the `tsmdbmgr.opt` file.
 - o *commethod* specifies the client API that is used to contact the server for database backup. This value can be `tcpip` or `sharedmem`. For more information about shared memory, see step 5.
 - o *tcpserveraddr* specifies the server address that the client API uses to contact the server for database backup. To ensure that the database can be backed up, this value must be `localhost`.
 - o *tcpport* specifies the port number that the client API uses to contact the server for database backup. Ensure that you enter the same `tcpport` value that is specified in the `dsmserv.opt` server options file.
 - o *errorlogname* specifies the error log where the client API logs errors that are encountered during a database backup. This log is typically in the server instance directory. However, this log can be placed in any location where the instance user ID has write-permission.
 - o *nodename* specifies the node name that the client API uses to connect to the server during a database backup. To ensure that the database can be backed up, this value must be `$_TSMDBMGR_`.
5. Optional: Configure the server to back up the database by using shared memory. In this way, you might be able to reduce the processor load and improve throughput. Complete the following steps:
- a. Review the `dsmserv.opt` file. If the following lines are not in the file, add them:

```
commethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

- b. In the `dsm.sys` configuration file, locate the following lines:

```
commethod tcpip
tcpserveraddr localhost
tcpport port_number
```

Replace the specified lines with the following lines:

```
commethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

AIX: Configuring server options for server database maintenance

To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.

About this task

Table and index reorganization requires significant processor resources, active log space, and archive log space. Because database backup takes precedence over reorganization, select the time and duration for reorganization to ensure that the processes do not overlap and reorganization can complete.

AIX You can optimize index and table reorganization for the server database. In this way, you can help to avoid unexpected database growth and performance issues. For instructions, see technote 1683633.

If you update these server options while the server is running, you must stop and restart the server before the updated values take effect.

Procedure

1. Modify the server options.

AIX Edit the server options file, `dsmserv.opt`, in the server instance directory. Follow these guidelines when you edit the server options file:

- o To enable an option, remove the asterisk at the beginning of the line.
- o Enter an option on any line.
- o Enter only one option per line. The entire option with its value must be on one line.
- o If you have multiple entries for an option in the file, the server uses the last entry.

To view available server options, see the sample file, `dsmserv.opt.smp`, in the `/opt/tivoli/tsm/server/bin` directory.

2. If you plan to use data deduplication, enable the `ALLOWREORGINDEX` server option. Add the following option and value to the server options file:

```
allowreorgindex yes
```

3. Set the REORGBEGINTIME and REORGDURATION server options to control when reorganization starts and how long it runs. Select a time and duration so that reorganization runs when you expect that the server is least busy. These server options control both table and index reorganization processes.
 - a. Set the time for reorganization to start by using the REORGBEGINTIME server option. Specify the time by using the 24-hour system. For example, to set the start time for reorganization as 8:30 p.m., specify the following option and value in the server options file:

```
reorgbegintime 20:30
```

- b. Set the interval during which the server can start reorganization. For example, to specify that the server can start reorganization for four hours after the time set by the REORGBEGINTIME server option, specify the following option and value in the server options file:

```
reorgduration 4
```

4. If the server was running while you updated the server options file, stop and restart the server.

AIX

AIX: Starting the server instance

You can start the server by using the instance user ID, which is the preferred method, or the root user ID.

Before you begin

Ensure that you set access permissions and user limits correctly.

AIX

For instructions, see [Verifying access rights and user limits](#).

About this task

When you start the server by using the instance user ID, you simplify the setup process and avoid potential issues. However, in some cases, it might be necessary to start the server with the root user ID. For example, you might want to use the root user ID to ensure that the server can access specific devices. You can set up the server to start automatically by using either the instance user ID or the root user ID.

AIX

If you must complete maintenance or reconfiguration tasks, start the server in maintenance mode.

Procedure

To start the server, take one of the following actions:

- Start the server by using the instance user ID.

AIX

For instructions, see [Starting the server from the instance user ID](#).

- Start the server by using the root user ID.

For instructions about authorizing root user IDs to start the server, see [Authorizing root user IDs to start the server \(V7.1.1\)](#). For instructions about starting the server by using the root user ID, see [Starting the server from the root user ID \(V7.1.1\)](#).

- **AIX** Start the server automatically.

AIX

For instructions, see [AIX: Automatically starting servers](#).

- **AIX** Start the server in maintenance mode.

For instructions, see [AIX: Starting the server in maintenance mode](#).

AIX

AIX: Verifying access rights and user limits

Before you start the server, verify access rights and user limits.

About this task

If you do not verify user limits, also known as *ulimits*, you might experience server instability or a failure of the server to respond. You must also verify the system-wide limit for the maximum number of open files. The system-wide limit must be greater than or equal to the user limit.

Procedure

1. Verify that the server instance user ID has permissions to start the server.
2. For the server instance that you plan to start, ensure that you have authority to read and write files in the server instance directory. Verify that the `dsmserv.opt` file exists in the server instance directory, and that the file includes parameters for the server instance.
3. If the server is attached to a tape drive, medium changer, or removable media device, and you plan to start the server by using the instance user ID, grant read/write access to the instance user ID for these devices. To set permissions, take one of the following actions:

- o If the system is dedicated to IBM Spectrum Protect™ and only the IBM Spectrum Protect administrator has access, make the device special file world-writable. On the operating system command line, issue the following command:

```
chmod +w /dev/rmtX
```

- o If the system has multiple users, you can restrict access by making the IBM Spectrum Protect instance user ID the owner of the special device files. On the operating system command line, issue the following command:

```
chmod u+w /dev/rmtX
```

- o If multiple user instances are running on the same system, change the group name, for example TAPEUSERS, and add each IBM Spectrum Protect instance user ID to that group. Then, change the ownership of the device special files to belong to the group TAPEUSERS, and make them group-writable. On the operating system command line, issue the following command:

```
chmod g+w /dev/rmtX
```

4. Verify the following user limits based on the guidelines in the table.

Table 1. User limit (ulimit) values

User limit type	Preferred value	Command to query value
Maximum size of core files created	Unlimited	<code>ulimit -Hc</code>
Maximum size of a data segment for a process	Unlimited	<code>ulimit -Hd</code>
Maximum file size	Unlimited	<code>ulimit -Hf</code>
Maximum number of open files	65536	<code>ulimit -Hn</code>
Maximum amount of processor time in seconds	Unlimited	<code>ulimit -Ht</code>

To modify user limits, follow the instructions in the documentation for your operating system.

Tip: If you plan to start the server automatically by using a script, you can set the user limits in the script.

5. Ensure that the user limit of maximum user processes (the `nproc` setting) is set to the minimum suggested value of 16384.
 - a. To verify the current user limit, issue the `ulimit -Hu` command by using the instance user ID. For example:

```
[user@Machine ~]$ ulimit -Hu
16384
```

- b. If the limit of maximum user processes is not set to 16384, set the value to 16384.

AIX Add the following line to the `/etc/security/limits` file:

```
instance_user_id          -          nproc          16384
```

where `instance_user_id` specifies the server instance user ID.

AIX

AIX: Starting the server from the instance user ID

To start the server from the instance user ID, log in with the instance user ID and issue the appropriate command from the server instance directory.

Before you begin

Ensure that access rights and user limits are set correctly. For instructions, see [AIX: Verifying access rights and user limits](#).

Procedure

1. Log in to the system where IBM Spectrum Protect™ is installed by using the instance user ID for the server.
2. If you do not have a user profile that runs the `db2profile` script, issue the following command:

```
. /home/tsminst1/sqlllib/db2profile
```

Tip: For instructions about updating the user ID login script to run the `db2profile` script automatically, see the [DB2® documentation](#).

3. Start the server by issuing the following command on one line from the server instance directory:

AIX

```
LDR_CNTRL=TEXTPSIZE=64K@DATAPSIZE=64K@STACKPSIZE=64K@SHMPSIZE=64K  
usr/bin/dmserv
```

AIX

Ensure that you include a space after `SHMPSIZE=64K`. By starting the server with this command, you enable 64 KB memory pages for the server. This setting helps you optimize server performance.

Tip: The command runs in the foreground so that you can set an administrator ID and connect to the server instance.

AIX

For example, if the name of the server instance is `tsminst1` and the server instance directory is `/tsminst1`, you can start the instance by issuing the following commands:

```
cd /tsminst1  
. ~/sqlllib/db2profile  
LDR_CNTRL=TEXTPSIZE=64K@DATAPSIZE=64K@STACKPSIZE=64K@SHMPSIZE=64K  
usr/bin/dmserv
```

AIX

AIX: Automatically starting servers

You can configure the server to start automatically at system startup. Use the `rc.dmserv` script, which is provided for this purpose.

Before you begin

Ensure that access rights and user limits are set correctly.

AIX

For instructions, see [Verifying access rights and user limits](#).

About this task

The `rc.dmserv` script is in the server installation directory, for example, in the `/opt/tivoli/tsm/server/bin` directory.

AIX

Tip: If you used the configuration wizard, you might have chosen to start the server automatically when the system is restarted. If you selected that choice, an entry for starting the server was added automatically to the `/etc/inittab` file.

Procedure

If you did not use a wizard to configure the server, add an entry to the `/etc/inittab` file for each server that you want to automatically start:

1. Set the run level to the value that corresponds to multiuser mode with networking enabled. Typically, the run level to use is 2, 3, or 5, depending on the operating system and its configuration. Ensure that the run level in the `/etc/inittab` file matches the run level of the operating system. For more information about multiuser mode and run levels, see the documentation for your operating system.

2. On the `rc.dsmsserv` command in the `/etc/inittab` file, specify the instance user ID with the `-u` option, and the location of the server instance directory with the `-i` option. If you want to start more than one server instance automatically, add an entry for each server instance. To verify the syntax, see the documentation for your operating system.
Tip: To automatically start a server instance with the root user ID, use the `-U` option.

Example

For example, if the instance owner is `tsminst1` and the server instance directory is `/home/tsminst1/tsminst1`, add the following entry to `/etc/inittab`, on one line:

AIX

```
tsm1:2:once:/opt/tivoli/tsm/server/bin/rc.dsmsserv -u tsminst1
-i /home/tsminst1/tsminst1 -q >/dev/console 2>&1
```

In this example, the ID for the process is `tsm1`, and the run level is set to 2.

If you have more than one server instance that you want to run, add an entry for each server instance. For example, if you have instance owner IDs `tsminst1` and `tsminst2`, and instance directories `/home/tsminst1/tsminst1` and `/home/tsminst2/tsminst2`, add the following entries to `/etc/inittab`. Each entry is on one line.

AIX

```
tsm1:2:once:/opt/tivoli/tsm/server/bin/rc.dsmsserv -u tsminst1
-i /home/tsminst1/tsminst1 -q >/dev/console 2>&1
tsm2:2:once:/opt/tivoli/tsm/server/bin/rc.dsmsserv -u tsminst2
-i /home/tsminst2/tsminst2 -q >/dev/console 2>&1
```

Related information:

[Server startup script: rc.dsmsserv](#)

AIX

AIX: Starting the server in maintenance mode

You can start the server in maintenance mode to avoid disruptions during maintenance and reconfiguration tasks.

About this task

Start the server in maintenance mode by running the `DSMSERV` utility with the `MAINTENANCE` parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmsserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmsserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode.

Operations that were disabled during maintenance mode are reenabled.

AIX: Stopping the server

You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.

About this task

To stop the server, issue the following command from the IBM Spectrum Protect™ command line:

```
halt
```

AIX If you cannot connect to the server with an administrative client and you want to stop the server, you must cancel the process by using the kill command with the process ID number (pid). The pid is displayed at initialization. **Important:** Before you issue the kill command, ensure that you know the correct process ID for the IBM Spectrum Protect server. The `dsmserve.v6lock` file, in the directory from which the server is running, can be used to identify the process ID of the process to kill. To display the file, enter:

```
cat /instance_dir/dsmserve.v6lock
```

AIX Issue the following command to stop the server:

```
kill -36 dsmserve_pid
```

where `dsmserve_pid` is the process ID number.

AIX: Registering licenses

Immediately register any IBM Spectrum Protect™ licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.

About this task

Use the REGISTER LICENSE command for this task. See REGISTER LICENSE for more details.

Example: Register a license

Register the base IBM Spectrum Protect license.

```
register license file=tsmbasic.lic
```

AIX: Preparing the server for database backup operations

To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.

Procedure

1. Ensure that the IBM Spectrum Protect™ configuration is complete. If you did not use the configuration wizard (`dsmicfgx`) to configure the server, ensure that you completed the steps to manually configure the server for database backups.
2. Select the device class to be used for database backups, protect the master encryption key, and set a password. All of these actions are completed by issuing the SET DBRECOVERY command from the administrative command line:

```
set dbrecovery device_class_name protectkeys=yes password=password_name
```

where *device_class_name* specifies the device class to be used for database backup operations, and *password_name* specifies the password.

You must specify a device class name or the backup fails. By specifying PROTECTKEYS=YES, you ensure that the master encryption key is backed up during database backup operations.

Important: Create a strong password that is at least 8 characters long. Ensure that you remember this password. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database.

Example

To specify that database backups include a copy of the master encryption key for the server, run the following command:

```
set dbrecovery dbback protectkeys=yes password=protect8991
```

AIX: Running multiple server instances on a single system

You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

Multiply the memory and other system requirements for one server by the number of instances planned for the system.

AIX The set of files for one instance of the server is stored separately from the files used by another server instance on the same system. Use the steps in AIX: Creating the server instance for each new instance, including creation of the new instance user.

To manage the system memory that is used by each server, use the DBMEMPERCENT server option to limit the percentage of system memory. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.

You can upgrade directly from V7.1 to V8.1. See the upgrade section (Upgrading to V8.1) for more details. When you upgrade and have multiple servers on your system, you must run the installation wizard only once. The installation wizard collects the database and variables information for all of your original server instances.

If you upgrade from IBM Spectrum Protect V6.3 to V8.1.5 and have multiple servers on your system, all instances that exist in DB2® V9.7 are dropped and recreated in DB2 V11.1. The wizard issues the `db2 upgrade db dbname` command for each database. The database environment variables for each instance on your system are also reconfigured during the upgrade process.

Related tasks:

[Running multiple server instances on a single system \(V7.1.1\)](#)

AIX: Monitoring the server

When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Procedure

1. Monitor the active log to ensure that the size is correct for the workload that is handled by the server instance.

When the server workload reaches its typical expected level, the space that is used by the active log is 80% - 90% of the space that is available to the active log directory. At that point, you might need to increase the amount of space. Whether you must increase the space depends on the types of transactions in the server workload. Transaction characteristics affect how the active log space is used.

The following transaction characteristics can affect the space usage in the active log:

- The number and size of files in backup operations
 - Clients such as file servers that back up large numbers of small files can cause large numbers of transactions that are completed quickly. The transactions might use a large amount of space in the active log, but for a short time.
 - Clients such as a mail server or a database server that back up large amounts of data in few transactions can cause small numbers of transactions that take a long time to complete. The transactions might use a small

- amount of space in the active log, but for a long time.
- Network connection types
 - Backup operations that occur over fast network connections cause transactions that complete more quickly. The transactions use space in the active log for a shorter time.
 - Backup operations that occur over relatively slower connections cause transactions that take a longer time to complete. The transactions use space in the active log for a longer time.

If the server is handling transactions with a wide variety of characteristics, the space that is used for the active log might increase and decrease significantly over time. For such a server, you might need to ensure that the active log typically has a smaller percentage of its space used. The extra space allows the active log to grow for transactions that take a long time to complete.

2. Monitor the archive log to ensure that space is always available.

Remember: If the archive log becomes full, and the failover archive log becomes full, the active log can become full, and the server stops. The goal is to make enough space available to the archive log so that it never uses all its available space.

You are likely to notice the following pattern:

- a. Initially, the archive log grows rapidly as typical client-backup operations occur.
- b. Database backups occur regularly, either as scheduled or done manually.
- c. After at least two full database backups occur, log pruning occurs automatically. The space that is used by the archive log decreases when the pruning occurs.
- d. Normal client operations continue, and the archive log grows again.
- e. Database backups occur regularly, and log pruning occurs as often as full database backups occur.

With this pattern, the archive log grows initially, decreases, and then might grow again. Over time, as normal operations continue, the amount of space that is used by the archive log should reach a relatively constant level.

If the archive log continues to grow, consider taking one or both of these actions:

- Add space to the archive log. You might need to move the archive log to a different file system.
 - Increase the frequency of full database backups, so that log pruning occurs more frequently.
3. If you defined a directory for the failover archive log, determine whether any logs get stored in that directory during normal operations. If the failover log space is being used, consider increasing the size of the archive log. The goal is that the failover archive log is used only under unusual conditions, not in normal operation.

AIX: Installing an IBM Spectrum Protect server fix pack

IBM Spectrum Protect™ maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.

Before you begin

To install a fix pack or interim fix to the server, install the server at the level on which you want to run it. You do not have to start the server installation at the base release level. For example, if you currently have V8.1.1 installed, you can go directly to the latest fix pack for V8.1. You do not have to start with the V8.1.0 installation if a maintenance update is available.

You must have the IBM Spectrum Protect license package installed. The license package is provided with the purchase of a base release. When you download a fix pack or interim fix from Fix Central, install the server license that is available on the Passport Advantage® website. To display messages and help in a language other than US English, install the language package of your choice.

If you upgrade the server to V8.1.5 or later, and then revert the server to a level that is earlier than V8.1.5, you must restore the database to a point in time before the upgrade. During the upgrade process, complete the required steps to ensure that the database can be restored: back up the database, the volume history file, the device configuration file, and the server options file. For more information, see AIX: Reverting from Version 8.1.5 to a previous server.

If you are using the client management service, ensure that you upgrade it to the same version as the IBM Spectrum Protect server.

Ensure that you retain the installation media from the base release of the installed server. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Visit the IBM® Support Portal for the following information:

- A list of the latest maintenance and download fixes. Click **Downloads** and apply any applicable fixes.
- Details about obtaining a base license package. Search for **Downloads > Passport Advantage**.
- Supported platforms and system requirements. Search for **IBM Spectrum Protect supported operating systems**.

Ensure that you upgrade the server before you upgrade backup-archive clients. If you do not upgrade the server first, communication between the server and clients might be interrupted.

Attention: Do not alter the DB2® software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Procedure

To install a fix pack or interim fix, complete the following steps:

1. Back up the database. The preferred method is to use a snapshot backup. A snapshot backup is a full database backup that does not interrupt any scheduled database backups. For example, issue the following IBM Spectrum Protect administrative command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information. Issue the following IBM Spectrum Protect administrative command:

```
backup devconfig filenames=file_name
```

where *file_name* specifies the name of the file in which to store device configuration information.

3. Save the volume history file to another directory or rename the file. Issue the following IBM Spectrum Protect administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named dmserv.opt. The file is in the server instance directory.
5. Halt the server before installing a fix pack or interim fix. Use the HALT command.
6. Ensure that extra space is available in the installation directory. The installation of this fix pack might require additional temporary disk space in the installation directory of the server. The amount of additional disk space can be as much as that required for installing a new database as part of an IBM Spectrum Protect installation. The IBM Spectrum Protect installation wizard displays the amount of space that is required for installing the fix pack and the available amount. If the required amount of space is greater than the available amount, the installation stops. If the installation stops, add the required disk space to the file system and restart the installation.
7. **AIX** Log in as the root user.
8. Obtain the package file for the fix pack or interim fix that you want to install from the IBM Support Portal, Passport Advantage, or Fix Central.
9. **AIX** Change to the directory where you placed the executable file and complete the following steps.

Tip: The files are extracted to the current directory. Ensure that the executable file is in the directory where you want the extracted files to be located.

- a. Change file permissions by entering the following command:

```
chmod a+x 8.x.x.x-IBM-SPSRV-platform.bin
```

where *platform* denotes the architecture that IBM Spectrum Protect is to be installed on.

- b. Issue the following command to extract the installation files:

```
./8.x.x.x-IBM-SPSRV-platform.bin
```

10. Select one of the following ways of installing IBM Spectrum Protect.

Important: After a fix pack is installed, it is not necessary to go through the configuration again. You can stop after completing the installation, fix any errors, then restart your servers.

Install the IBM Spectrum Protect software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

AIX: Installing IBM Spectrum Protect by using the installation wizard

Tip: After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.

Command line in console mode

Follow the instructions for your operating system:
AIX: Installing IBM Spectrum Protect by using console mode

Silent mode

Follow the instructions for your operating system:
AIX: Installing IBM Spectrum Protect in silent mode

Tip: If you have multiple server instances on your system, run the installation wizard only once. The installation wizard upgrades all server instances.

Results

Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click File > View Log. To collect log files, from the IBM Installation Manager tool, click Help > Export Data for Problem Analysis.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

- `AIX /var/ibm/InstallationManager/logs`
- `AIX AIX: Applying a fix pack to IBM Spectrum Protect V8.1.5 in a clustered environment for AIX`
IBM Spectrum Protect maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level. It is possible to apply a fix pack onto a clustered environment for AIX®.

AIX: Reverting from Version 8.1.5 to a previous server

If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect™ server with minimal loss of data.

Before you begin

You must have the following items from the earlier version of the server:

- Server database backup
- Volume history file
- Device configuration file
- Server options file

About this task

Use the same instructions whether you are reverting within releases or to an earlier release, for example, from 8.1.3 to 8.1.2 or from 8.1.3 to 7.1.2. The older version must match the version that you used before the upgrade to V8.1.

Attention: Specify the REUSEDELAY parameter to help prevent backup-archive client data loss when you revert the server to a previous version.

Steps for reverting to the previous server version

About this task

Complete the following steps on the system that has the V8.1 server.

Procedure

1. Halt the server to shut down all server operations by using the HALT command.
2. Remove the database from the database manager, then delete the database and recovery log directories.
 - a. Manually remove the database. One way to remove it is by issuing this command: `AIX`

```
dsmserv removedb tsmdb1
```

b. If you must reuse the space that is occupied by the database and recovery log directories, you can now delete these directories.

3. Use the uninstallation program to uninstall the V8.1 server. Uninstallation removes the server and the database manager, with their directories. For details, see AIX: Uninstalling IBM Spectrum Protect.
4. Stop the cluster service. Reinstall the version of the server program that you were using before the upgrade to V8.1.5. This version must match the version that your server was running when you created the database backup that you restore in a later step. For example, the server was at V7.1.7 before the upgrade, and you intend to use the database backup that was in use on this server. You must install the V7.1.7 fix pack to be able to restore the database backup.
5. Configure the new server database by using the configuration wizard. To start the wizard, issue the following command:

```
AIX  
./dsmicfgx
```

6. Ensure that no servers are running in the background.
7. Restore the database to a point in time before the upgrade.
8. Copy the following files to the instance directory.
 - o Device configuration file
 - o Volume history file
 - o The server options file (typically dsmserv.opt)
9. If you enabled data deduplication for any FILE-type storage pools that existed before the upgrade, or if you moved data that existed before the upgrade into new storage pools while using the V8.1.5 server, you must complete additional recovery steps. For more details, see Additional recovery steps if you created new storage pools or enabled data deduplication.
10. If the REUSEDDELAY parameter setting on storage pools is less than the age of the database that you restored, restore volumes on any sequential-access storage pools that were reclaimed after that database backup. Use the RESTORE VOLUME command.
If you do not have a backup of a storage pool, audit the reclaimed volumes by using the AUDIT VOLUME command, with the FIX=YES parameter to resolve inconsistencies. For example:

```
audit volume volume_name fix=yes
```
11. If client backup or archive operations were completed using the V8.1 server, audit the storage pool volumes on which the data was stored.

Additional recovery steps if you created new storage pools or enabled data deduplication

If you created new storage pools, turned on data deduplication for any FILE-type storage pools, or did both while your server was running as a V8.1.5 server, you must complete more steps to return to the previous server version.

Before you begin

To complete this task, you must have a complete backup of the storage pool that was created before the upgrade to V8.1.5.

About this task

Use this information if you did either or both of the following actions while your server was running as a V8.1.5 server:

- You enabled the data deduplication function for any storage pools that existed before the upgrade to V8.1.5 program. Data deduplication applies only to storage pools that use a FILE device type.
- You created new primary storage pools after the upgrade *and* moved data that was stored in other storage pools into the new storage pools.

Complete these steps after the server is again restored to V7.

Procedure

- For each storage pool for which you enabled the data deduplication function, restore the entire storage pool by using the RESTORE STGPOOL command.
- For storage pools that you created after the upgrade, determine what action to take. Data that was moved from existing V8 storage pools into the new storage pools might be lost because the new storage pools no longer exist in your restored V8 server. Possible recovery depends on the type of storage pool:
 - o If data was moved from V8 DISK-type storage pools into a new storage pool, space that was occupied by the data that was moved was probably reused. Therefore, you must restore the original V8 storage pools by using the storage

pool backups that were created before the upgrade to V8.1.5.

If no data was moved from V8 DISK-type storage pools into a new storage pool, then audit the storage pool volumes in these DISK-type storage pools.

- o If data was moved from V8 sequential-access storage pools into a new storage pool, that data might still exist and be usable in storage pool volumes on the restored V8 server. The data might be usable if the REUSEDELAY parameter for the storage pool was set to a value that prevented reclamation while the server was running as a V8.1.5 server. If any volumes were reclaimed while the server was running as a V8.1.5 server, restore those volumes from storage pool backups that were created before the upgrade to V8.1.5.

AIX: Reference: DB2 commands for IBM Spectrum Protect server databases

Use this list as reference when you are directed to issue DB2® commands by IBM® support.

Purpose

After using the wizards to install and configure IBM Spectrum Protect™, you seldom need to issue DB2 commands. A limited set of DB2 commands that you might use or be asked to issue are listed in Table 1. This list is supplemental material only and is not a comprehensive list. There is no implication that an IBM Spectrum Protect administrator will use it on a daily or ongoing basis. Samples of some commands are provided. Details of output are not listed.

For a full explanation of the commands described here and of their syntax, see the DB2 product information.

Table 1. DB2 commands

Command	Description	Example
db2icrt	<p>Creates DB2 instances in the home directory of the instance owner.</p> <p>Tip: The IBM Spectrum Protect configuration wizard creates the instance used by the server and database. After a server is installed and configured through the configuration wizard, the db2icrt command is generally not used.</p> <p>AIX This utility is in the DB2DIR/instance directory, where DB2DIR represents the installation location where the current version of the DB2 database system is installed.</p>	<p>Manually create an IBM Spectrum Protect instance. Enter the command on one line:</p> <pre>/opt/tivoli /tsm/db2/in stance/ db2icrt -a server -u instance_na me instance_na me</pre>
db2set	Displays DB2 variables.	<p>List DB2 variables:</p> <pre>db2set</pre>
CATALOG DATABASE	Stores database location information in the system database directory. The database can be located either on the local workstation or on a remote database partition server. The server configuration wizard takes care of any catalog needed for using the server database. Run this command manually, after a server is configured and running, only if something in the environment changes or is damaged.	<p>Catalog the database:</p> <pre>db2 catalog database tsmdb1</pre>

Command	Description	Example
CONNECT TO DATABASE	Connects to a specified database for command-line interface (CLI) use.	Connect to the IBM Spectrum Protect database from a DB2 CLI: <pre>db2 connect to tsmdbl</pre>
GET DATABASE CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Show the configuration information for a database alias: <pre>db2 get db cfg for tsmdbl</pre> Retrieve information in order to verify settings such as database configuration, log mode, and maintenance. <pre>db2 get db config for tsmdbl show detail</pre>
GET DATABASE MANAGER CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Retrieve configuration information for the database manager: <pre>db2 get dbm cfg</pre>
GET HEALTH SNAPSHOT	Retrieves the health status information for the database manager and its databases. The information returned represents a snapshot of the health state at the time the command was issued. IBM Spectrum Protect monitors the state of the database using the health snapshot and other mechanisms that are provided by DB2. There might be cases where the health snapshot or other DB2 documentation indicates that an item or database resource might be in an alert state. Such a case indicates that action must be considered to remedy the situation. IBM Spectrum Protect monitors the condition and responds appropriately. Not all declared alerts by the DB2 database are acted on.	Receive a report on DB2 health monitor indicators: <pre>db2 get health snapshot for database on tsmdbl</pre>

Command	Description	Example
GRANT (Database Authorities)	Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.	Grant access to the user ID itmuser: db2 GRANT CONNECT ON DATABASE TO USER itmuser db2 GRANT CREATETAB ON DATABASE TO USER itmuser
RUNSTATS	Updates statistics about the characteristics of a table and associated indexes or statistical views. These characteristics include number of records, number of pages, and average record length. To see a table, issue this utility after updating or reorganizing the table. A view must be enabled for optimization before its statistics can be used to optimize a query. A view that is enabled for optimization is known as a statistical view. Use the DB2 ALTER VIEW statement to enable a view for optimization. Issue the RUNSTATS utility when changes to underlying tables substantially affect the rows returned by the view. Tip: The server configures DB2 to run the RUNSTATS command as needed.	Update statistics on a single table. db2 runstats on table SCHEMA_NAME .TABLE_NAME with distribution and sampled detailed indexes all
SET SCHEMA	Changes the value of the CURRENT SCHEMA special register, in preparation for issuing SQL commands directly through the DB2 CLI. Tip: A special register is a storage area that is defined for an application process by the database manager. It is used to store information that can be referenced in SQL statements.	Set the schema for IBM Spectrum Protect: db2 set schema tsmdb1
START DATABASE MANAGER	Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts. Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Start the database manager: db2start
STOP DATABASE MANAGER	Stops the current database manager instance. Unless explicitly stopped, the database manager continues to be active. This command does not stop the database manager instance if any applications are connected to databases. If there are no database connections, but there are instance attachments, the command forces the instance attachments to stop first. Then, it stops the database manager. This command also deactivates any outstanding database activations before stopping the database manager. This command is not valid on a client. The server starts and stops the instance and database whenever the server starts and halts. Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Stop the database manager: db2 stop dbm

AIX: Uninstalling IBM Spectrum Protect

You can use the following procedures to uninstall IBM Spectrum Protect™. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Before you begin

Complete the following steps before you uninstall IBM Spectrum Protect:

- Complete a full database backup.
- Save a copy of the volume history and device configuration files.
- Store the output volumes in a safe location.

About this task

You can uninstall IBM Spectrum Protect by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

- **AIX: Uninstalling IBM Spectrum Protect by using a graphical wizard**
You can uninstall IBM Spectrum Protect by using the IBM® Installation Manager installation wizard.
- **AIX: Uninstalling IBM Spectrum Protect in console mode**
To uninstall IBM Spectrum Protect by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.
- **AIX: Uninstalling IBM Spectrum Protect in silent mode**
To uninstall IBM Spectrum Protect in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.
- **AIX: Uninstalling and reinstalling IBM Spectrum Protect**
If you plan to manually reinstall IBM Spectrum Protect instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.
- **AIX: Uninstalling IBM Installation Manager**
You can uninstall IBM Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

What to do next

See AIX: Installing the server components for installation steps to reinstall the IBM Spectrum Protect components.

AIX: Uninstalling IBM Spectrum Protect by using a graphical wizard

You can uninstall IBM Spectrum Protect™ by using the IBM® Installation Manager installation wizard.

Procedure

1. Start the Installation Manager.

AIX In the directory where the Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command:

```
./IBMIM
```

2. Click Uninstall.
3. Select IBM Spectrum Protect server, and click Next.
4. Click Uninstall.
5. Click Finish.

AIX: Uninstalling IBM Spectrum Protect in console mode

To uninstall IBM Spectrum Protect™ by using the command line, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **AIX** eclipse/tools

For example:

- o `AIX` /opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command:
 - o `AIX` `./imcl -c`
 3. To uninstall, enter 5.
 4. Choose to uninstall from the IBM Spectrum Protect package group.
 5. Enter N for Next.
 6. Choose to uninstall the IBM Spectrum Protect server package.
 7. Enter N for Next.
 8. Enter U for Uninstall.
 9. Enter F for Finish.

AIX: Uninstalling IBM Spectrum Protect in silent mode

To uninstall IBM Spectrum Protect™ in silent mode, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the IBM Spectrum Protect server components. IBM Spectrum Protect includes a sample response file, `uninstall_response_sample.xml`, in the input directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

If you want to uninstall all IBM Spectrum Protect components, leave `modify="false"` set for each component in the response file. If you do not want to uninstall a component, set the value to `modify="true"`.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o `AIX` `eclipse/tools`

For example:

- o `AIX` /opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command, where `response_file` represents the response file path, including the file name:

```
AIX
./imcl -input response_file -silent
```

The following command is an example:

```
AIX
./imcl -input /tmp/input/uninstall_response.xml -silent
```

AIX: Uninstalling and reinstalling IBM Spectrum Protect

If you plan to manually reinstall IBM Spectrum Protect™ instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.

About this task

To manually uninstall and reinstall IBM Spectrum Protect, complete the following steps:

1. `AIX` Make a list of your current server instances before proceeding to the uninstallation. Run the following command:

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

2. Run the following commands for every server instance:

```
AIX
```

```
db2 attach to instance_name
db2 get dbm cfg show detail
db2 detach
```

Keep a record of the database path for each instance.

3. Uninstall IBM Spectrum Protect. See AIX: Uninstalling IBM Spectrum Protect.
4. When you uninstall any supported version of IBM Spectrum Protect, including a fix pack, an instance file is created. The instance file is created to help reinstall IBM Spectrum Protect. Check this file and use the information when you are prompted for the instance credentials when reinstalling. In silent installation mode, you provide these credentials using the `INSTANCE_CRED` variable.

You can find the instance file in the following location:

- o `AIX` /etc/tivoli/tsm/instanceList.obj

5. Reinstall IBM Spectrum Protect. See AIX: Installing the server components.

If the `instanceList.obj` file does not exist, you need to recreate your server instances using the following steps:

- a. Recreate your server instances. See AIX: Creating the server instance.

Tip: The installation wizard configures the server instances but you must verify that they exist. If they do not exist, you must manually configure them.

- b. Catalog the database. Log in to each server instance as the instance user, one at a time, and issue the following commands:

`AIX`

```
db2 catalog database tsmdb1
db2 attach to instance_name
db2 update dbm cfg using dftdbpath instance_directory
db2 detach
```

- c. `AIX` Verify that the server instance was created successfully. Issue this command:

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

- d. Verify that IBM Spectrum Protect recognizes the server instance by listing your directories. Your home directory appears if you did not change it. Your instance directory does appear if you used the configuration wizard. Issue this command:

```
db2 list database directory
```

If you see TSMDB1 listed, you can start the server.

AIX: Uninstalling IBM Installation Manager

You can uninstall IBM® Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

Before you begin

Before you uninstall IBM Installation Manager, you must ensure that all packages that were installed by IBM Installation Manager are uninstalled. Close IBM Installation Manager before you start the uninstall process.

`AIX` To view installed packages, issue the following command from a command line:

```
cd /opt/IBM/InstallationManager/eclipse/tools
./imcl listInstalledPackages
```

Procedure

To uninstall IBM Installation Manager, complete the following steps:

`AIX`

1. Open a command line and change directories to `/var/ibm/InstallationManager/uninstall`.
2. Issue the following command:

```
./uninstall
```

Restriction: You must be logged in to the system as the `root` user ID.

Linux: Installing the server

Installation of the server includes planning, installation, and initial configuration.

- **Linux: Planning to install the server**
Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.
- **Linux: Installing the server components**
To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.
- **Linux: Taking the first steps after you install IBM Spectrum Protect**
After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect instance.
- **Linux: Installing an IBM Spectrum Protect server fix pack**
IBM Spectrum Protect maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.
- **Linux: Reverting from Version 8.1.5 to a previous server**
If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect server with minimal loss of data.
- **Linux: Reference: DB2 commands for IBM Spectrum Protect server databases**
Use this list as reference when you are directed to issue DB2® commands by IBM® support.
- **Linux: Uninstalling IBM Spectrum Protect**
You can use the following procedures to uninstall IBM Spectrum Protect. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Linux: Planning to install the server

Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.

- **Linux: What you should know first**
Before installing IBM Spectrum Protect, be familiar with your operating systems, storage devices, communication protocols, and system configurations.
- **Linux: Planning for optimal performance**
Before you install the IBM Spectrum Protect server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.
- **Linux** **Linux: Minimum system requirements for Linux systems**
To install the IBM Spectrum Protect server on a Linux system, it is necessary to have a minimum level of hardware and software, including a communication method and the most current device driver.
- **Linux** **Linux: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system**
You can install other products that deploy and use DB2® products on the same system as the IBM Spectrum Protect Version 8.1.5 server, with some limitations.
- **Linux: IBM Installation Manager**
IBM Spectrum Protect uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.
- **Linux: Worksheets for planning details for the server**
You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect server. You can also use them to keep track of names and user IDs.
- **Linux: Capacity planning**
Capacity planning for IBM Spectrum Protect includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.
- **Linux: Server naming best practices**
Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect server.
- **Linux: Installation directories**
Installation directories for the IBM Spectrum Protect server include the server, DB2, device, language, and other directories. Each one contains several additional directories.

Linux: What you should know first

Before installing IBM Spectrum Protect™, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

Server maintenance releases, client software, and publications are available from the IBM® Support Portal.

Linux Restriction: You can install and run the Version 8.1.5 server on a system that already has DB2® installed on it, whether DB2 was installed independently or as part of some other application, with some restrictions. For details, see the compatibility with other DB2 products topic.

Experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. Do not, however, use DB2 tools to change DB2 configuration settings from those that are preset by IBM Spectrum Protect, or alter the DB2 environment for IBM Spectrum Protect in other ways, such as with other products. The V8.1.5 server has been built and tested extensively using the data definition language (DDL) and database configuration that the server deploys.

Attention: Do not alter the DB2 software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Linux: Planning for optimal performance

Before you install the IBM Spectrum Protect™ server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.

Procedure

1. Review Linux: What you should know first.
2. Review each of the following sub-sections.
 - Linux: Planning for the server hardware and the operating system
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - Linux: Planning for the server database disks
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - Linux: Planning for the server recovery log disks
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
 - Linux: Planning for directory-container and cloud-container storage pools
Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.
 - Linux: Planning for storage pools in DISK or FILE device classes
Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.
 - Linux: Planning for the correct type of storage technology
Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect.
 - Linux: Applying best practices to the server installation
Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Linux: Planning for the server hardware and the operating system

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
<p>Does the operating system and hardware meet or exceed requirements?</p> <ul style="list-style-type: none"> • Number and speed of processors • System memory • Supported operating system level 	<p>If you are using the minimum required amount of memory, you can support a minimal workload.</p> <p>You can experiment by adding more system memory to determine whether the performance is improved. Then, decide whether you want to keep the system memory dedicated to the server. Test the memory variations by using the entire daily cycle of the server workload.</p> <p>If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.</p>	<p>Review operating system requirements at technote 1243309.</p> <p>Additionally, review the guidance in Tuning tasks for operating systems and other applications.</p> <p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For more information about sizing requirements for the server and storage, see the IBM Spectrum Protect™ Blueprint.</p>
<p>Are disks configured for optimal performance?</p>	<p>The amount of tuning that can be done for different disk systems varies. Ensure that the appropriate queue depths and other disk system options are set.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"

Question	Tasks, characteristics, options, or settings	More information
Does the server have enough memory?	<p>Heavier workloads and advanced features such as data deduplication and node replication require more than the minimum system memory that is specified in the system requirements document.</p> <p>For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements:</p> <ul style="list-style-type: none"> • For databases less than 500 GB, you need 16 GB of memory. • For databases with a size of 500 GB - 1 TB, you need 24 GB of memory. • For databases with a size of 1 TB - 1.5 TB, you need 32 GB of memory. • For databases greater than 1.5 TB, you need 40 GB of memory. <p>Ensure that you allocate extra space for the active log and the archive log for replication processing.</p>	<p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication • Memory requirements
Does the system have enough host bus adapters (HBAs) to handle the data operations that the IBM Spectrum Protect server must run simultaneously?	<p>Understand what operations require use of HBAs at the same time.</p> <p>For example, a server must store 1 GB/sec of backup data while also doing storage pool migration that requires 0.5 GB/sec capacity to complete. The HBAs must be able to handle all of the data at the speed required.</p>	See Tuning HBA capacity.

Question	Tasks, characteristics, options, or settings	More information
Is network bandwidth greater than the planned maximum throughput for backups?	<p>Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments.</p> <p>For node replication, network bandwidth must be greater than the planned maximum throughput.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Tuning network performance • Checklist for node replication
Are you using a preferred file system for IBM Spectrum Protect server files?	<p>Use a file system that ensures optimal performance and data availability. The server uses direct I/O with file systems that support the feature. Using direct I/O can improve throughput and reduce processor use. For more information about the preferred file system for your operating system, see IBM Spectrum Protect server-supported file systems.</p>	<p>For more information, see Configuring the operating system for disk performance.</p>
Are you planning to configure enough paging space?	<p>Paging space, or swap space, extends the memory that is available for processing. When the amount of free RAM in the system is low, programs or data that is not in use are moved from memory to paging space. This action releases memory for other activities, such as database operations.</p> <p>Linux Use a minimum of 32 GB of paging space or 50% of your RAM, whichever value is larger.</p>	
Linux Are you planning to tune the kernel parameters after installation of the server?	Linux You must tune kernel parameters.	Linux See the information about tuning kernel parameters: Linux: Tuning kernel parameters for Linux systems

Linux: Planning for the server database disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
Is the database on fast, low-latency disks?	<p>Do not use the following drives for the IBM Spectrum Protect™ database:</p> <ul style="list-style-type: none"> • Nearline SAS (NL-SAS) • Serial Advanced Technology Attachment (SATA) • Parallel Advanced Technology Attachment (PATA) <p>Do not use internal disks that are included by default in most server hardware.</p> <p>Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.</p> <p>If you plan to use the data deduplication functions of IBM Spectrum Protect, focus on disk performance in terms of I/O operations per second (IOPS).</p>	For more information, see Checklist for data deduplication.
Is the database stored on disks or LUNs that are separate from disks or LUNs that are used for the active log, archive log, and storage pool volumes?	<p>Separation of the server database from other server components helps reduce contention for the same resources by different operations that must run at the same time.</p> <p>Tip: The database and the archive log can share an array when you use solid-state drive (SSD) technology.</p>	
If you are using RAID, do you know how to select the optimal RAID level for your system? Are you defining all LUNs with the same size and type of RAID?	<p>When a system must do large numbers of writes, RAID 10 outperforms RAID 5. However, RAID 10 requires more disks than RAID 5 for the same amount of usable storage.</p> <p>If your disk system is RAID, define all your LUNs with the same size and type of RAID. For example, do not mix 4+1 RAID 5 with 4+2 RAID 6.</p>	
If an option to set the strip size or segment size is available, are you planning to optimize the size when you configure the disk system?	If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.	The block size that is used for the database varies depending on the table space. Most table spaces use 8 KB blocks, but some use 32 KB blocks.

Question	Tasks, characteristics, options, or settings	More information
<p>Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database?</p> <p>Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.</p>	<p>Heavier workloads and use of some features require more database storage paths than the minimum requirements.</p> <p>Server operations such as data deduplication drive a high number of input/output operations per second (IOPS) for the database. Such operations perform better when the database has more directories.</p> <p>For server databases that are larger than 2 TB or are expected to grow to that size, use eight directories.</p> <p>Consider planned growth of the system when you determine how many storage paths to create. The server uses the higher number of storage paths more effectively if the storage paths are present when the server is first created.</p> <p>Use the <i>DB2_PARALLEL_IO</i> variable to force parallel I/O to occur on table spaces that have one container, or on table spaces that have containers on more than one physical disk. If you do not set the <i>DB2_PARALLEL_IO</i> variable, I/O parallelism is equal to the number of containers that are used by the table space. For example, if a table space spans four containers, the level of I/O parallelism that is used is 4.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For help with forecasting growth when the server deduplicates data, see technote 1596944.</p> <p>For the most recent information about database size, database reorganization, and performance considerations for IBM Spectrum Protect servers, see technote 1683633.</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>
<p>Are all directories for the database the same size?</p>	<p>Directories that are all the same size ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching.</p> <p>This guideline also applies if you must add storage paths after the initial configuration of the server.</p>	
<p>Are you planning to raise the queue depth of the database LUNs on AIX® systems?</p>	<p>The default queue depth is often too low.</p>	<p>See Configuring AIX systems for disk performance.</p>

Linux: Planning for the server recovery log disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Are the active log and archive log stored on disks or LUNs that are separate from what is used for the database and storage pool volumes?</p>	<p>Ensure that the disks where you place the active log are not used for other server or system purposes. Do not place the active log on disks that contain the server database, the archive log, or system files such as page or swap space.</p>	<p>Separation of the server database, active log, and archive log helps to reduce contention for the same resources by different operations that must run at the same time.</p>

Question	Tasks, characteristics, options, or settings	More information
Are the logs on disks that have nonvolatile write cache?	Nonvolatile write cache allows data to be written to the logs as fast as possible. Faster write operations for the logs can improve performance for server operations.	
Are you setting the logs to a size that adequately supports the workload?	<p>If you are not sure about the workload, use the largest size that you can.</p> <p>Active log The maximum size is 512 GB, set with the ACTIVELOGSIZE server option.</p> <p>Ensure that there is at least 8 GB of free space on the active log file system after the fixed size active logs are created.</p> <p>Archive log The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Make the archive log at least as large as the active log.</p>	<ul style="list-style-type: none"> For log sizing details, see the recovery log information in technote 1421060. For information about sizing when you use data deduplication, see Checklist for data deduplication.
Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?	The archive failover log is for emergency use by the server when the archive log becomes full. Slower disks can be used for the archive failover log.	<p>Use the ARCHFAILOVERLOGDIRECTORY server option to specify the location of the archive failover log.</p> <p>Monitor the usage of the directory for the archive failover log. If the archive failover log must be used by the server, the space for the archive log might not be large enough.</p>
If you are mirroring the active log, are you using only one type of mirroring?	<p>You can mirror the log by using one of the following methods. Use only one type of mirroring for the log.</p> <ul style="list-style-type: none"> Use the MIRRORLOGDIRECTORY option that is available for the IBM Spectrum Protect™ server to specify a mirror location. Use software mirroring, such as Logical Volume Manager (LVM) on AIX®. Use mirroring in the disk system hardware. 	<p>If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability.</p> <p>For more information, see Configuring and tuning the recovery log.</p>

Linux: Planning for directory-container and cloud-container storage pools

Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
<p>Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the IBM Spectrum Protect™ database?</p>	<p>Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing.</p> <p>Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication), add 1000 IOPS to this minimum.</p> <p>For example, an IBM Spectrum Protect server that is ingesting 3 TB of data per day would need 6000 IOPS for the database disks:</p> $3000 \text{ IOPS minimum} + 3000 (3 \text{ TB} \times 1000 \text{ IOPS}) = 6000 \text{ IOPS}$	<p>For recommendations about disk selection, see "Planning for server database disks".</p> <p>For more information about IOPS, see the IBM Spectrum Protect Blueprints.</p>
<p>Do you have enough memory for the size of your database?</p>	<p>Use a minimum of 40 GB of system memory for IBM Spectrum Protect servers, with a database size of 100 GB, that are deduplicating data. If the retained capacity of backup data grows, the memory requirement might need to be higher.</p> <p>Monitor memory usage regularly to determine whether more memory is required.</p> <p>Use more system memory to improve caching of database pages. The following memory size guidelines are based on the daily amount of new data that you back up:</p> <ul style="list-style-type: none"> • 128 GB of system memory for daily backups of data, where the database size is 1 - 2 TB • 192 GB of system memory for daily backups of data, where the database size is 2 - 4 TB 	<p>Memory requirements</p>

Question	Tasks, characteristics, options, or settings	More information
<p>Have you properly sized the storage capacity for the database active log and archive log?</p>	<p>Configure the server to have a minimum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072.</p> <p>The suggested starting size for the archive log is 1 TB. The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Ensure that there is at least 10% extra disk space for the file system than the size of the archive log.</p> <p>Use a directory for the database archive logs with an initial free capacity of at least 1 TB. Specify the directory by using the ARCHLOGDIRECTORY server option.</p> <p>Define space for the archive failover log by using the ARCHFAILOVERLOGDIRECTORY server option.</p>	<p>For more information about sizing for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Is compression enabled for the archive log and database backups?</p>	<p>Enable the ARCHLOGCOMPRESS server option to save storage space.</p> <p>This compression option is different from inline compression. Inline compression is enabled by default with IBM Spectrum Protect V7.1.5 and later.</p> <p>Restriction: Do not use this option if the amount of backed up data exceeds 6 TB per day.</p>	<p>For more information about compression for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Are the IBM Spectrum Protect database and logs on separate disk volumes (LUNs)?</p> <p>Is the disk that is used for the database configured according to best practices for a transactional database?</p>	<p>The database must not share disk volumes with IBM Spectrum Protect database logs or storage pools, or with any other application or file system.</p>	<p>For more information about server database and recovery log configuration, see Server database and recovery log configuration and tuning.</p>
<p>Are you using a minimum of eight (2.2 GHz or equivalent) processor cores for each IBM Spectrum Protect server that you plan to use with data deduplication?</p>	<p>If you are planning to use client-side data deduplication, verify that client systems have adequate resources available during a backup operation to complete data deduplication processing. Use a processor that is at least the minimum equivalent of one 2.2 GHz processor core per backup process with client-side data deduplication.</p>	<ul style="list-style-type: none"> • Effective planning and use of deduplication • IBM Spectrum Protect Blueprints

Question	Tasks, characteristics, options, or settings	More information
Did you allocate enough storage space for the database?	<p>For a rough estimate, plan for 100 GB of database storage for every 50 TB of data that is to be protected in deduplicated storage pools. <i>Protected data</i> is the amount of data before data deduplication, including all versions of objects stored.</p> <p>As a best practice, define a new container storage pool exclusively for data deduplication. Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.</p>	
Have you estimated storage pool capacity to configure enough space for the size of your environment?	<p>You can estimate capacity requirements for a deduplicated storage pool by using the following technique:</p> <ol style="list-style-type: none"> 1. Estimate the base size of the source data. 2. Estimate the daily backup size by using an estimated change and growth rate. 3. Determine retention requirements. 4. Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements. 5. Apply the deduplication ratio factor. 6. Apply the compression ratio factor. 7. Round up the estimate to consider transient storage pool usage. 	For an example of using this technique, see Effective planning and use of deduplication.
Have you distributed disk I/O over many disk devices and controllers?	<p>Use arrays that consist of as many disks as possible, which is sometimes referred to as wide striping. Ensure that you use one database directory per distinct array on the subsystem.</p> <p>Set the <i>DB2_PARALLEL_IO</i> registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks.</p> <p>When I/O bandwidth is available and the files are large, for example 1 MB, the process of finding duplicates can occupy the resources of an entire processor. When files are smaller, other bottlenecks can occur.</p> <p>Specify eight or more file systems for the deduplicated storage pool device class so that I/O is distributed across as many LUNs and physical devices as possible.</p>	<p>For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes".</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>

Question	Tasks, characteristics, options, or settings	More information
Have you scheduled daily operations based on your backup strategy?	<p>The best practice sequence of operations is in the following order:</p> <ol style="list-style-type: none"> 1. Client backup 2. Storage pool protection 3. Node replication 4. Database backup 5. Expire inventory 	<ul style="list-style-type: none"> • Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools
Do you have enough storage to manage the DB2® lock list?	<p>If you deduplicate data that includes large files or large numbers of files concurrently, the process can result in insufficient storage space. When the lock list storage is insufficient, backup failures, data management process failures, or server outages can occur.</p> <p>File sizes greater than 500 GB that are processed by data deduplication are most likely to deplete storage space. However, if many backup operations use client-side data deduplication, this problem can also occur with smaller-sized files.</p>	For information about tuning the DB2 LOCKLIST parameter, see Tuning server-side data deduplication.
Is sufficient bandwidth available to transfer data to an IBM Spectrum Protect server?	<p>To transfer data to an IBM Spectrum Protect server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required.</p> <p>Use a V7.1.5 server or higher to use inline compression and use a V7.1.6 or later client to enable enhanced compression processing.</p>	For more information, see the enablededup client option.
Have you determined how many storage pool directories to assign to each storage pool?	<p>Assign directories to a storage pool by using the DEFINE STGPOOLDIRECTORY command.</p> <p>Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).</p>	

Question	Tasks, characteristics, options, or settings	More information
<p>Did you allocate enough disk space in the cloud-container storage pool?</p>	<p>To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space:</p> <ul style="list-style-type: none"> • For serial-attached SCSI (SAS) and spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount, in terabytes, for disk space. • Provide 3 TB for flash-based storage systems with fast network connections to on-premises, high-performance cloud systems. • Provide 5 TB for solid-state drive (SSD) systems with fast network connections to high-performance cloud systems. 	
<p>Did you select the appropriate type of local storage?</p>	<p>Ensure that data transfers from local storage to cloud finish before the next backup cycle starts. Tip: Data is removed from local storage soon after it moves to the cloud. Use the following guidelines:</p> <ul style="list-style-type: none"> • Use flash or SSD for large systems that have high-performing cloud systems. Ensure that you have a dedicated 10 GB wide area network (WAN) link with a high-speed connection to the object storage. For example, use flash or SSD if you have a dedicated 10 GB WAN link plus a high-speed connection to either an IBM® Cloud Object Storage location or to an Amazon Simple Storage Service (Amazon S3) data center. • Use larger capacity 15000 rpm SAS disks for these scenarios: <ul style="list-style-type: none"> ◦ Medium-sized systems ◦ Slower cloud connections, for example, 1 GB ◦ When you use IBM Cloud Object Storage as your service provider across several regions • For SAS or spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount for disk space, in terabytes. 	

Linux: Planning for storage pools in DISK or FILE device classes

Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.

Question	Tasks, characteristics, options, or settings	More information
Can the storage pool LUNs sustain throughput rates for 256 KB sequential reads and writes to adequately handle the workload within the time constraints?	<p>When you are planning for peak loads, consider all the data that you want the server to read or write to the disk storage pools simultaneously. For example, consider the peak flow of data from client backup operations and server data-movement operations such as migration that run at the same time.</p> <p>The IBM Spectrum Protect™ server reads and writes to storage pools predominantly in 256 KB blocks.</p> <p>If the disk system includes the capability, configure the disk system for optimal performance with sequential read/write operations rather than random read/write operations.</p>	For more information, see Analyzing the basic performance of disk systems.
Is the disk configured to use read and write cache?	Use more cache for better performance.	
For storage pools that use FILE device classes, have you determined a good size to use for the storage pool volumes?	Review the information in Optimal number and size of volumes for storage pools that use disk. If you do not have the information to estimate a size for FILE device class volumes, start with volumes that are 50 GB.	Typically, problems arise more frequently when the volumes are too small. Few problems are reported when volumes are larger than needed. When you determine the volume size to use, as a precaution choose a size that might be larger than necessary.
For storage pools that use FILE device classes, are you using preallocated volumes?	<p>Scratch volumes can cause file fragmentation.</p> <p>To ensure that a storage pool does not run out of volumes, set the MAXSCRATCH parameter to a value greater than zero.</p>	<p>Use the DEFINE VOLUME server command to preallocate volumes in the storage pool.</p> <p>Use the DEFINE STGPOOL or UPDATE STGPOOL server command to set the MAXSCRATCH parameter.</p>
For storage pools that use FILE device classes, have you compared the maximum number of client sessions to the number of volumes that are defined?	Always maintain enough usable volumes in the storage pools to allow for the expected peak number of client sessions that run at one time. The volumes might be scratch volumes, empty volumes, or partly filled volumes.	For storage pools that use FILE device classes, only one session or process can write to a volume at the same time.

Question	Tasks, characteristics, options, or settings	More information
<p>For storage pools that use FILE device classes, have you set the MOUNTLIMIT parameter of the device class to a value that is high enough to account for the number of volumes that might be mounted in parallel?</p>	<p>For storage pools that use data deduplication, the MOUNTLIMIT parameter is typically in the range of 500 - 1000.</p> <p>Set the value for MOUNTLIMIT to the maximum number of mount points that are needed for all active sessions. Consider parameters that affect the maximum number of mount points that are needed:</p> <ul style="list-style-type: none"> • The MAXSESSIONS server option, which is the maximum number of IBM Spectrum Protect sessions that can run concurrently. • The MAXNUMMP parameter, which sets the maximum number of mount points that each client node can use. <p>For example, if the maximum number of client node backup sessions is typically 100 and each of the nodes has MAXNUMMP=2, multiply 100 nodes by the 2 mount points for each node to get the value of 200 for the MOUNTLIMIT parameter.</p>	<p>Use the REGISTER NODE or UPDATE NODE server command to set the MAXNUMMP parameter for client nodes.</p>
<p>For storage pools that use DISK device classes, have you determined how many storage pool volumes to put on each file system?</p>	<p>How you configure the storage for a storage pool that uses a DISK device class depends on whether you are using RAID for the disk system.</p> <p>If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system.</p> <p>If you are using RAID 5 with $n + 1$ volumes, configure the storage in one of the following ways:</p> <ul style="list-style-type: none"> • Configure n file systems on the LUN and define one storage pool volume per file system. • Configure one file system and n storage pool volumes for the LUN. 	<p>For an example layout that follows this guideline, see Sample layout of server storage pools.</p>
<p>Did you create your storage pools to distribute I/O across multiple file systems?</p>	<p>Ensure that each file system is on a different LUN on the disk system.</p> <p>Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.</p>	<p>For details, see the following topics:</p> <ul style="list-style-type: none"> • Tuning disk storage for the server • Tuning and configuring storage pools and volumes

Linux: Planning for the correct type of storage technology

Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect™.

Procedure

Review the following table to help you to choose the correct type of storage technology for the storage resources that the server requires.

Table 1. Storage technology types for IBM Spectrum Protect storage requirements

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
Solid-state disk (SSD)	Place the database on SSD in the following circumstances: <ul style="list-style-type: none"> You are using IBM Spectrum Protect data deduplication. You are backing up more than 8 TB of new data daily. 	If you place the IBM Spectrum Protect database on an SSD, as a best practice, place the active log on an SSD. If space is not available, use high-performance disk instead.	Save SSDs for use with the database and active log. The archive log and archive failover logs can be placed on slower storage technology types.	Save SSDs for use with the database and active log. Storage pools can be placed on slower storage technology types.
High-performance disk with the following characteristics: <ul style="list-style-type: none"> 15k rpm disk Fibre Channel or serial-attached SCSI (SAS) interface 	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. Isolate the server database from its logs and storage pools, and from data for other applications.	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use high-performance disks for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use high-performance disks for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.
Medium-performance or high-performance disk with the following characteristics: <ul style="list-style-type: none"> 10k rpm disk Fibre Channel or SAS interface 	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. Isolate the server database from its logs and storage pools, and from data for other applications.	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use medium-performance or high-performance disk for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use medium-performance or high-performance disk for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
SATA, network-attached storage	Do not use this storage for the database. Do not place the database on XIV storage systems.	Do not use this storage for the active log.	Use of this slower storage technology is acceptable because these logs are written once and infrequently read.	Use this slower storage technology in the following circumstances: <ul style="list-style-type: none"> • Data is infrequently written, for example written once. • Data is infrequently read.
Tape and virtual tape				Use for long-term retention or if data is infrequently used.

Linux: Applying best practices to the server installation

Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect™ solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Procedure

- The following best practices are the most important for optimal performance and problem prevention.
- Review the table to determine the best practices that apply to your environment.

Best practice	More information
Use fast disks for the server database. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.	Use fast, low-latency disks for the database. Using SSD is essential if you are using data deduplication and node replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: <ul style="list-style-type: none"> ◦ "Planning for server database disks" ◦ "Planning for the correct type of storage technology"
Ensure that the server system has enough memory.	Review operating system requirements in technote 1243309. Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and node replication can require more than the minimum memory that is specified in the system requirements document. If you plan to run multiple instances, each instance requires the memory that is listed for one server. Multiply the memory for one server by the number of instances that are planned for the system.

Best practice	More information
Separate the server database, the active log, the archive log, and disk storage pools from each other.	<p>Keep all IBM Spectrum Protect storage resources on separate disks. Keep storage pool disks separate from the disks for the server database and logs. Storage pool operations can interfere with database operations when both are on the same disks. Ideally, the server database and logs are also separated from each other. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> o "Planning for server database disks" o "Planning for server recovery log disks" o "Planning for storage pools in DISK or FILE device classes"
Use at least four directories for the server database. For larger servers or servers that use advanced features, use eight directories.	<p>Place each directory on a LUN that is isolated from other LUNs and from other applications.</p> <p>A server is considered to be large if its database is larger than 2 TB or is expected to grow to that size. Use eight directories for such servers.</p> <p>See "Planning for server database disks".</p>
If you are using data deduplication, node replication, or both, follow the guidelines for database configuration and other items.	<p>Configure the server database according to the guidelines, because the database is extremely important to how well the server runs when these features are being used. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> o Checklist for data deduplication o Checklist for node replication
For storage pools that use FILE type device classes, follow the guidelines for the size of storage pool volumes. Typically, 50 GB volumes are best.	<p>Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size.</p> <p>Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements.</p> <p>Isolate the storage devices that are used by IBM Spectrum Protect from other applications that have high I/O, and ensure that there is enough throughput to that storage.</p> <p>For more details, see Checklist for storage pools on DISK or FILE.</p>
Schedule IBM Spectrum Protect client operations and server maintenance activities to avoid or minimize overlap of operations.	<p>For more details, see the following topics:</p> <ul style="list-style-type: none"> o Tuning the schedule for daily operations o Checklist for server configuration
Monitor operations constantly.	<p>By monitoring, you can find problems early and more easily identify causes. Keep records of monitoring reports for up to a year to help you identify trends and plan for growth. See Monitoring and maintaining the environment for performance.</p>

Linux: Minimum system requirements for Linux systems

To install the IBM Spectrum Protect™ server on a Linux system, it is necessary to have a minimum level of hardware and software, including a communication method and the most current device driver.

The optimal IBM Spectrum Protect environment is set up with data deduplication by using the IBM Spectrum Protect Blueprints.

The IBM Spectrum Protect device driver package does not contain a device driver for this operating system because a SCSI generic device driver is used. Configure the device driver before using the IBM Spectrum Protect server with tape devices. The IBM Spectrum Protect driver package contains driver tools and ACSLS daemons. You can locate IBM® driver packages at the Fix Central website.

Requirements, supported devices, client installation packages, and fixes are available in the IBM Support Portal for IBM Spectrum Protect. After you install IBM Spectrum Protect and before you customize it for your use, go to the website and download and apply any applicable fixes.

- **Linux** Linux: Minimum Linux x86_64 server requirements
Before you install an IBM Spectrum Protect server on a Linux x86_64 operating system, review the hardware and software requirements.
- **Linux** Linux: Minimum Linux on System z server requirements
Before you install an IBM Spectrum Protect server on a Linux on System z® operating system, review the hardware and software requirements.
- **Linux** Linux: Minimum Linux on Power Systems (little endian) server requirements
Before you install an IBM Spectrum Protect server on a Linux on Power Systems (little endian) operating system, review the hardware and software requirements.

Linux: Minimum Linux x86_64 server requirements

Before you install an IBM Spectrum Protect™ server on a Linux x86_64 operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Spectrum Protect server installation

For the most current information about IBM Spectrum Protect system requirements, see technote 1243309.

Linux: Minimum Linux on System z server requirements

Before you install an IBM Spectrum Protect™ server on a Linux on System z® operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Spectrum Protect server installation

For the most current information about IBM Spectrum Protect system requirements, see technote 1243309.

Linux: Minimum Linux on Power Systems™ (little endian) server requirements

Before you install an IBM Spectrum Protect™ server on a Linux on Power Systems (little endian) operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Spectrum Protect server installation

For the most current information about IBM Spectrum Protect system requirements, see technote 1243309.

Linux

Linux: Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

You can install other products that deploy and use DB2® products on the same system as the IBM Spectrum Protect™ Version 8.1.5 server, with some limitations.

To install and use other products that use a DB2 product on the same system as the IBM Spectrum Protect server, ensure that the following criteria are met:

Table 1. Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

Cri ter ion	Instructions
-------------------	--------------

Criterion	Instructions
Version level	The other products that use a DB2 product must use DB2 version 9 or later. DB2 products include product encapsulation and segregation support beginning with Version 9. Starting with this version, you can run multiple copies of DB2 products, at different code levels, on the same system. For details, see the information about multiple DB2 copies in the DB2 product information.
User IDs and directories	Ensure that the user IDs, fence user IDs, installation location, other directories, and related information are not shared across DB2 installations. Your specifications must be different from the IDs and locations that you used for the IBM Spectrum Protect server installation and configuration. If you used the ds micfgx wizard to configure the server, these are values that you entered when running the wizard. If you used the manual configuration method, review the procedures that you used if necessary to recall the values that were used for the server.
Resource allocation	<p>Consider the resources and capability of the system compared to the requirements for both the IBM Spectrum Protect server and the other applications that use the DB2 product. To provide sufficient resources for the other DB2 applications, you might have to change the IBM Spectrum Protect server settings so that the server uses less system memory and resources. Similarly, if the workloads for the other DB2 applications compete with the IBM Spectrum Protect server for processor or memory resources, the performance of the server in handling the expected client workload or other server operations might be adversely affected.</p> <p>To segregate resources and provide more capability for the tuning and allocation of processor, memory, and other system resources for multiple applications, consider using logical partition (LPAR), workload partition (WPAR), or other virtual workstation support. For example, run a DB2 application on its own virtualized system.</p>

Linux: IBM Installation Manager

IBM Spectrum Protect™ uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install IBM Spectrum Protect. It must remain installed on the system so that IBM Spectrum Protect can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The IBM Spectrum Protect offering contains all of the media that IBM Installation Manager requires to install IBM Spectrum Protect.

Package

The group of software components that are required to install an offering.

The IBM Spectrum Protect package contains the following components:

- IBM Installation Manager installation program
- IBM Spectrum Protect offering

Package group

A set of packages that share a common parent directory.

The default package group for the IBM Spectrum Protect package is `IBM Installation Manager`.

Repository

A remote or local storage area for data and other application resources.

The IBM Spectrum Protect package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of IBM Spectrum Protect.

Linux: Worksheets for planning details for the server

You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect™ server. You can also use them to keep track of names and user IDs.

Item	Space required	Number of directories	Location of directories
The database			
Active log			
Archive log			
Optional: Log mirror for the active log			
Optional: Secondary archive log (failover location for archive log)			

Item	Names and user IDs	Location
The <i>instance user ID</i> for the server, which is the ID you use to start and run the IBM Spectrum Protect server		
The <i>home directory</i> for the server, which is the directory that contains the instance user ID		
The database instance name		
The <i>instance directory</i> for the server, which is a directory that contains files specifically for this server instance (the server options file and other server-specific files)		
The server name, use a unique name for each server		

Linux: Capacity planning

Capacity planning for IBM Spectrum Protect™ includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.

- **Linux: Estimating space requirements for the database**
To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.
- **Linux: Recovery log space requirements**
In IBM Spectrum Protect, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.
- **Linux: Monitoring space utilization for the database and recovery logs**
To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

- Linux: Deleting installation rollback files
You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

Linux: Estimating space requirements for the database

To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.

About this task

Consider using at least 25 GB for the initial database space. Provision file system space appropriately. A database size of 25 GB is adequate for a test environment or a library-manager-only environment. For a production server supporting client workloads, the database size is expected to be larger. If you use random-access disk (DISK) storage pools, more database and log storage space is needed than for sequential-access storage pools.

The maximum size of the IBM Spectrum Protect™ database is 6 TB.

For information about sizing the database in a production environment that is based on the number of files and on storage pool size, see the following topics.

- Linux: Estimating database space requirements based on the number of files
If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.
- Linux: Estimating database space requirements based on storage pool capacity
To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.
- Linux: The database manager and temporary space
The IBM Spectrum Protect server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

Linux: Estimating database space requirements based on the number of files

If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

About this task

To estimate space requirements that is based on the maximum number of files in server storage, use the following guidelines:

- 600 - 1000 bytes for each stored version of a file, including image backups.
Restriction: The guideline does not include space that is used during data deduplication.
- 100 - 200 bytes for each cached file, copy storage pool file, active-data pool file, and deduplicated file.
- Additional space is required for database optimization to support varying data-access patterns and to support server back-end processing of the data. The amount of extra space is equal to 50% of the estimate for the total number of bytes for file objects.

In the following example for a single client, the calculations are based on the maximum values in the preceding guidelines. The examples do not take into account that you might use file aggregation. In general, when you aggregate small files, it reduces the amount of required database space. File aggregation does not affect space-managed files.

Procedure

1. Calculate the number of file versions. Add each of the following values to obtain the number of file versions:
 - a. Calculate the number of backed-up files. For example, as many as 500,000 client files might be backed up at a time.
In this example, storage policies are set to keep up to three copies of backed up files:

$$500,000 \text{ files} * 3 \text{ copies} = 1,500,000 \text{ files}$$

- b. Calculate the number of archive files. For example, as many as 100,000 client files might be archived copies.
- c. Calculate the number of space-managed files. For example, as many as 200,000 client files might be migrated from client workstations.

Using 1000 bytes per file, the total amount of database space that is required for the files that belong to the client is 1.8 GB:

$$(1,500,000 + 100,000 + 200,000) * 1000 = 1.8 \text{ GB}$$

2. Calculate the number of cached files, copy storage-pool files, active-data pool files, and deduplicated files:
 - a. Calculate the number of cached copies. For example, caching is enabled in a 5 GB disk storage pool. The high migration threshold of the pool is 90% and the low migration threshold of the pool is 70%. Thus, 20% of the disk pool, or 1 GB, is occupied by cached files.
If the average file size is about 10 KB, approximately 100,000 files are in cache at any one time:

$$100,000 \text{ files} * 200 \text{ bytes} = 19 \text{ MB}$$

- b. Calculate the number of copy storage-pool files. All primary storage pools are backed up to the copy storage pool:

$$(1,500,000 + 100,000 + 200,000) * 200 \text{ bytes} = 343 \text{ MB}$$

- c. Calculate the number of active storage-pool files. All the active client-backup data in primary storage pools is copied to the active-data storage pool. Assume that 500,000 versions of the 1,500,000 backup files in the primary storage pool are active:

$$500,000 * 200 \text{ bytes} = 95 \text{ MB}$$

- d. Calculate the number of deduplicated files. Assume that a deduplicated storage pool contains 50,000 files:

$$50,000 * 200 \text{ bytes} = 10 \text{ MB}$$

Based on the preceding calculations, about 0.5 GB of extra database space is required for the client's cached files, copy storage-pool files, active-data pool files, and deduplicated files.

3. Calculate the amount of extra space that is required for database optimization. To provide optimal data access and management by the server, extra database space is required. The amount of extra database space is equal to 50% of the total space requirements for file objects.

$$(1.8 + 0.5) * 50\% = 1.2 \text{ GB}$$

4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

$$1.8 + 0.5 + 1.2 = 3.5 \text{ GB}$$

5. Calculate the total amount of database space that is required for all clients. If the client that was used in the preceding calculations is typical and you have 500 clients, for example, you can use the following calculation to estimate the total amount of database space that is required for all clients:

$$500 * 3.5 = 1.7 \text{ TB}$$

Results

Tip: In the preceding examples, the results are estimates. The actual size of the database might differ from the estimate because of factors such as the number of directories and the length of the path and file names. Periodically monitor your database and adjust its size as necessary.

What to do next

During normal operations, the IBM Spectrum Protect™ server might require temporary database space. This space is needed for the following reasons:

- To hold the results of sorting or ordering that are not already being kept and optimized in the database directly. The results are temporarily held in the database for processing.
- To give administrative access to the database through one of the following methods:
 - A DB2® open database connectivity (ODBC) client
 - An Oracle Java™ database connectivity (JDBC) client
 - Structured Query Language (SQL) to the server from an administrative-client command line

Consider using an extra 50 GB of temporary space for every 500 GB of space for file objects and optimization. See the guidelines in the following table. In the example that is used in the preceding step, a total of 1.7 TB of database space is required for file objects and optimization for 500 clients. Based on that calculation, 200 GB is required for temporary space. The total amount of required database space is 1.9 TB.

Database size	Minimum temporary-space requirement
< 500 GB	50 GB
≥ 500 GB and < 1 TB	100 GB
≥ 1 TB and < 1.5 TB	150 GB
≥ 1.5 and < 2 TB	200 GB
≥ 2 and < 3 TB	250 - 300 GB
≥ 3 and < 4 TB	350 - 400 GB

Linux: Estimating database space requirements based on storage pool capacity

To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.

Linux: The database manager and temporary space

The IBM Spectrum Protect™ server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

The database manager sorts data in a specific sequence, according to the SQL statement that you issue to request the data. Depending on the workload on the server, and if there is more data than the database manager can manage, the data (that is ordered in sequence) is allocated to temporary disk space. Data is allocated to temporary disk space when there is a large result set. The database manager dynamically manages the memory that is used when data is allocated to temporary disk space.

For example, expiration processing can produce a large result set. If there is not enough system memory on the database to store the result set, some of the data is allocated to temporary disk space. During expiration processing, if a node or file space are selected that are too large to process, the database manager cannot sort the data in memory. The database manager must use temporary space to sort data.

To run database operations, consider adding more database space for the following scenarios:

- The database has a small amount of space and the server operation that requires temporary space uses the remaining free space.
- The file spaces are large, or the file spaces have an assigned policy that creates many file versions.
- The IBM Spectrum Protect server must run with limited memory. The database uses the IBM Spectrum Protect server main memory to run database operations. However, if there is insufficient memory available, the IBM Spectrum Protect server allocates temporary space on disk to the database. For example, if 10G of memory is available and database operations require 12G of memory, the database uses temporary space.
- An `out of database space` error is displayed when you deploy an IBM Spectrum Protect server. Monitor the server activity log for messages that are related to database space.

Important: Do not change the DB2 software that is installed with the IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack, of DB2 software to avoid damage to the database.

Linux: Recovery log space requirements

In IBM Spectrum Protect™, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.

- Linux: Active and archive log space
When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.
- Linux: Active-log mirror space
The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.
- Linux: Archive-failover log space
The archive failover log is used by the server if the archive log directory runs out of space.

Linux: Active and archive log space

When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.

In IBM Spectrum Protect™ servers V7.1 and later, the active log can be a maximum size of 512 GB. The archive log size is limited to the size of the file system that it is installed on.

Use the following general guidelines when you estimate the size of the active log:

- The suggested starting size for the active log is 16 GB.
- Ensure that the active log is at least large enough for the amount of concurrent activity that the server typically handles. As a precaution, try to anticipate the largest amount of work that the server manages at one time. Provision the active log with extra space that can be used if needed. Consider using 20% of extra space.
- Monitor used and available active log space. Adjust the size of the active log as needed, depending upon factors such as client activity and the level of server operations.
- Ensure that the directory that holds the active log is as large as, or larger than, the size of the active log. A directory that is larger than the active log can accommodate failovers, if they occur.
- Ensure that the file system that contains the active log directory has at least 8 GB of free space for temporary log movement requirements.

The suggested starting size for the archive log is 48 GB.

The archive log directory must be large enough to contain the log files that are generated since the previous full backup. For example, if you perform a full backup of the database every day, the archive log directory must be large enough to hold the log files for all the client activity that occurs during 24 hours. To recover space, the server deletes obsolete archive log files after a full backup of the database. If the archive log directory becomes full and a directory for archive failover logs does not exist, log files remain in the active log directory. This condition can cause the active log directory to fill up and stop the server. When the server restarts, some of the existing active-log space is released.

After the server is installed, you can monitor archive log utilization and the space in the archive log directory. If the space in the archive log directory fills up, it can cause the following problems:

- The server is unable to perform full database backups. Investigate and resolve this problem.
- Other applications write to the archive log directory, exhausting the space that is required by the archive log. Do not share archive log space with other applications including other IBM Spectrum Protect servers. Ensure that each server has a separate storage location that is owned and managed by that specific server.
- Linux: Example: Estimating active and archive log sizes for basic client-store operations
Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.
- Linux: Example: Estimating active and archive log sizes for clients that use multiple sessions
If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.
- Linux: Example: Estimating active and archive log sizes for simultaneous write operations
If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.
- Linux: Example: Estimating active and archive log sizes for basic client store operations and server operations
Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.
- Linux: Example: Estimating active and archive log sizes under conditions of extreme variation
Problems with running out of active log space can occur if you have many transactions that complete quickly and some

transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

- Linux: Example: Estimating archive log sizes with full database backups
The IBM Spectrum Protect server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.
- Linux: Example: Estimating active and archive log sizes for data deduplication operations
If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

Linux: Example: Estimating active and archive log sizes for basic client-store operations

Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.

To determine the sizes of the active and archive logs for basic client-store operations, use the following calculation:

```
number of clients x files stored during each transaction
x log space needed for each file
```

This calculation is used in the example in the following table.

Table 1. Basic client-store operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053 bytes	The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	19.5 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 bytes = 3.5 GB Increase that amount by the suggested starting size of 16 GB: 3.5 + 16 = 19.5 GB
Archive log: Suggested size	58.5 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. 3.5 x 3 = 10.5 GB Increase that amount by the suggested starting size of 48 GB: 10.5 + 48 = 58.5 GB

Item	Example values	Description
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Linux: Example: Estimating active and archive log sizes for clients that use multiple sessions

If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.

To determine the sizes of the active and archive logs when clients use multiple sessions, use the following calculation:

$$\text{number of clients} \times \text{sessions for each client} \times \text{files stored during each transaction} \times \text{log space needed for each file}$$

This calculation is used in the example in the following table.

Table 1. Multiple client sessions

Item	Example values		Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	1000	The number of client nodes that back up, archive, or migrate files every night.
Possible sessions for each client	3	3	The setting of the client option RESOURCEUTILIZATION is larger than the default. Each client session runs a maximum of three sessions in parallel.
Files stored during each transaction	4096	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053	3053	<p>The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>

Item	Example values		Description
Active log: Suggested size	26.5 GB ¹	51 GB ¹	<p>The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 3 sessions for each client x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 10.5 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>10.5 + 16 = 26.5 GB</p> <p>The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.</p> <p>(1000 clients x 3 sessions for each client x 4096 files store during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 35 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>35 + 16 = 51 GB</p>
Archive log: Suggested size	79.5 GB ¹	153 GB ¹	<p>Because of the requirement to be able to store archive logs across three server-database backup cycles, the estimate for the active log is multiplied by 3:</p> <p>10.5 x 3 = 31.5 GB</p> <p>35 x 3 = 105 GB</p> <p>Increase those amounts by the suggested starting size of 48 GB:</p> <p>31.5 + 48 = 79.5 GB</p> <p>105 + 48 = 153 GB</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your active log and adjust its size if necessary.</p>			

Linux: Example: Estimating active and archive log sizes for simultaneous write operations

If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.

The log space that is required for each file increases by about 200 bytes for each copy storage pool that is used for a simultaneous write operation. In the example in the following table, data is stored to two copy storage pools in addition to a primary storage pool. The estimated log size increases by 400 bytes for each file. If you use the suggested value of 3053 bytes of log space for each file, the total number of required bytes is 3453.

This calculation is used in the example in the following table.

Table 1. Simultaneous write operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.

Item	Example values	Description
Log space that is required for each file	3453 bytes	3053 bytes plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3453 bytes for each file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size	60 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Linux: Example: Estimating active and archive log sizes for basic client store operations and server operations

Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.

For example, migration of files from the random-access (DISK) storage pool to a sequential-access disk (FILE) storage pool uses approximately 110 bytes of log space for each file that is migrated. For example, suppose that you have 300 backup-archive clients and each one of them backs up 100,000 files every night. The files are initially stored on DISK and then migrated to a FILE storage pool. To estimate the amount of active log space that is required for the data migration, use the following calculation. The number of clients in the calculation represents the maximum number of client nodes that back up, archive, or migrate files concurrently at any time.

300 clients x 100,000 files for each client x 110 bytes = 3.1 GB

Add this value to the estimate for the size of the active log that calculated for basic client store operations.

Linux: Example: Estimating active and archive log sizes under conditions of extreme variation

Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

Linux: Example: Estimating archive log sizes with full database backups

The IBM Spectrum Protect™ server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.

For example, if a full database backup occurs once a week, the archive log space must be able to contain the information in the archive log for a full week.

The difference in archive log size for daily and full database backups is shown in the example in the following table.

Table 1. Full database backups

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3453 bytes	3053 bytes for each file plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files per transaction x 3453 bytes per file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size with a full database backup every day	60 GB ¹	Because of the requirement to be able to store archive logs across three backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB

Item	Example values	Description
Archive log: Suggested size with a full database every week	132 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. Multiply the result by the number of days between full database backups:</p> $(4 \text{ GB} \times 3) \times 7 = 84 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $84 + 48 = 132 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested starting size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Linux: Example: Estimating active and archive log sizes for data deduplication operations

If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

The following factors affect requirements for active and archive log space:

The amount of deduplicated data

The effect of data deduplication on the active log and archive log space depends on the percentage of data that is eligible for deduplication. If the percentage of data that can be deduplicated is relatively high, more log space is required.

The size and number of extents

Approximately 1,500 bytes of active log space are required for each extent that is identified by a duplicate-identification process. For example, if 250,000 extents are identified by a duplicate-identification process, the estimated size of the active log is 358 MB:

$$250,000 \text{ extents identified during each process} \times 1,500 \text{ bytes for each extent} = 358 \text{ MB}$$

Consider the following scenario. Three hundred backup-archive clients back up 100,000 files each night. This activity creates a workload of 30,000,000 files. The average number of extents for each file is two. Therefore, the total number of extents is 60,000,000, and the space requirement for the archive log is 84 GB:

$$60,000,000 \text{ extents} \times 1,500 \text{ bytes for each extent} = 84 \text{ GB}$$

A duplicate-identification process operates on aggregates of files. An aggregate consists of files that are stored in a given transaction, as specified by the TXNGROUPMAX server option. Suppose that the TXNGROUPMAX server option is set to the default of 4096. If the average number of extents for each file is two, the total number of extents in each aggregate is 8192, and the space required for the active log is 12 MB:

$$8192 \text{ extents in each aggregate} \times 1500 \text{ bytes for each extent} = 12 \text{ MB}$$

The timing and number of the duplicate-identification processes

The timing and number of duplicate-identification processes also affects the size of the active log. Using the 12 MB active-log size that was calculated in the preceding example, the concurrent load on the active log is 120 MB if 10 duplicate-identification processes are running in parallel:

$$12 \text{ MB for each process} \times 10 \text{ processes} = 120 \text{ MB}$$

File size

Large files that are processed for duplicate identification can also affect the size of the active log. For example, suppose that a backup-archive client backs up an 80 GB, file-system image. This object can have a high number of duplicate extents if, for example, the files included in the file system image were backed up incrementally. For example, assume that a file

system image has 1.2 million duplicate extents. The 1.2 million extents in this large file represent a single transaction for a duplicate-identification process. The total space in the active log that is required for this single object is 1.7 GB:

$$1,200,000 \text{ extents} \times 1,500 \text{ bytes for each extent} = 1.7 \text{ GB}$$

If other, smaller duplicate-identification processes occur at the same time as the duplicate-identification process for a single large object, the active log might not have enough space. For example, suppose that a storage pool is enabled for deduplication. The storage pool has a mixture of data, including many relatively small files that range from 10 KB to several hundred KB. The storage pool also has few large objects that have a high percentage of duplicate extents.

To take into account not only space requirements but also the timing and duration of concurrent transactions, increase the estimated size of the active log by a factor of two. For example, suppose that your calculations for space requirements are 25 GB (23.3 GB + 1.7 GB for deduplication of a large object). If deduplication processes are running concurrently, the suggested size of the active log is 50 GB. The suggested size of the archive log is 150 GB.

The examples in the following tables show calculations for active and archive logs. The example in the first table uses an average size of 700 KB for extents. The example in the second table uses an average size of 256 KB. As the examples show, the average deduplicate-extent size of 256 KB indicates a larger estimated size for the active log. To minimize or prevent operational problems for the server, use 256 KB to estimate the size of the active log in your production environment.

Table 1. Average duplicate-extent size of 700 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	700 KB	700 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average size for extents.
Extents for a given file	1,198,372 bits	6,135,667 bits	Using the average extent size (700 KB), these calculations represent the total number of extents for a given object. The following calculation was used for an 800 GB object: $(800 \text{ GB} \div 700 \text{ KB}) = 1,198,372 \text{ bits}$ The following calculation was used for a 4 TB object: $(4 \text{ TB} \div 700 \text{ KB}) = 6,135,667 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	1.7 GB	8.6 GB	The estimated active log space that are needed for this transaction.

Item	Example values		Description
Active log: Suggested total size	66 GB ¹	79.8 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of two. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 1.7 \text{ GB}) \times 2 = 50 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $50 + 16 = 66 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 8.6 \text{ GB}) \times 2 = 63.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $63.8 + 16 = 79.8 \text{ GB}$
Archive log: Suggested size	198 GB ¹	239.4 GB ¹	<p>Multiply the estimated size of the active log by a factor of 3.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $50 \text{ GB} \times 3 = 150 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $150 + 48 = 198 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $63.8 \text{ GB} \times 3 = 191.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $191.4 + 48 = 239.4 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Table 2. Average duplicate-extent size of 256 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.

Item	Example values		Description
Average size of extents	256 KB	256 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average extent size.
Extents for a given file	3,276,800 bits	16,777,216 bits	<p>Using the average extent size, these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(800 \text{ GB} \div 256 \text{ KB}) = 3,276,800 \text{ bits}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(4 \text{ TB} \div 256 \text{ KB}) = 16,777,216 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	4.5 GB	23.4 GB	The estimated size of the active log space that is required for this transaction.
Active log: Suggested total size	71.6 GB ¹	109.4 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of 2. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 4.5 \text{ GB}) \times 2 = 55.6 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $55.6 + 16 = 71.6 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 23.4 \text{ GB}) \times 2 = 93.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $93.4 + 16 = 109.4 \text{ GB}$

Item	Example values		Description
Archive log: Suggested size	214.8 GB ¹	328.2 GB ¹	<p>The estimated size of the active log multiplied by a factor of 3.</p> <p>The following calculation was used for an 800 GB object:</p> $55.6 \text{ GB} \times 3 = 166.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $166.8 + 48 = 214.8 \text{ GB}$ <p>The following calculation was used for a 4 TB object:</p> $93.4 \text{ GB} \times 3 = 280.2 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $280.2 + 48 = 328.2 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Linux: Active-log mirror space

The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.

Creating a log mirror is a suggested option. If you increase the size of the active log, the log mirror size is increased automatically. Mirroring the log can affect performance because of the doubled I/O activity that is required to maintain the mirror. The additional space that the log mirror requires is another factor to consider when deciding whether to create a log mirror.

If the mirror log directory becomes full, the server issues error messages to the activity log and to the db2diag.log. Server activity continues.

Linux: Archive-failover log space

The archive failover log is used by the server if the archive log directory runs out of space.

Specifying an archive failover log directory can prevent problems that occur if the archive log runs out of space. If both the archive log directory and the drive or file system where the archive failover log directory is located become full, the data remains in the active log directory. This condition can cause the active log to fill up, which causes the server to halt.

Linux: Monitoring space utilization for the database and recovery logs

To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

Active log

If the amount of available active log space is too low, the following messages are displayed in the activity log:

ANR4531I: IC_AUTOBACKUP_LOG_USED_SINCE_LAST_BACKUP_TRIGGER

This message is displayed when the active log space exceeds the maximum specified size. The IBM Spectrum Protect™ server starts a full database backup.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR0297I: IC_BACKUP_NEEDED_LOG_USED_SINCE_LAST_BACKUP

This message is displayed when the active log space exceeds the maximum specified size. You must back up the database manually.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR4529I: IC_AUTOBACKUP_LOG_UTILIZATION_TRIGGER

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. If at least one full database backup has occurred, the IBM Spectrum Protect server starts an incremental database backup. Otherwise, the server starts a full database backup.

ANR0295I: IC_BACKUP_NEEDED_LOG_UTILIZATION

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. You must back up the database manually.

Archive log

If the amount of available archive log space is too low, the following message is displayed in the activity log:

ANR0299I: IC_BACKUP_NEEDED_ARCHLOG_USED

The ratio of used archive-log space to available archive-log space exceeds the log utilization threshold. The IBM Spectrum Protect server starts a full automatic database backup.

Database

If the amount of space available for database activities is too low, the following messages are displayed in the activity log:

ANR2992W: IC_LOG_FILE_SYSTEM_UTILIZATION_WARNING_2

The used database space exceeds the threshold for database space utilization. To increase the space for the database, use the EXTEND DBSPACE command, the EXTEND DBSPACE command, or the DSMSERV FORMAT utility with the DBDIR parameter.

ANR1546W: FILESYSTEM_DBPATH_LESS_1GB

The available space in the directory where the server database files are located is less than 1 GB.

When an IBM Spectrum Protect server is created with the DSMSERV FORMAT utility or with the configuration wizard, a server database and recovery log are also created. In addition, files are created to hold database information used by the database manager. The path specified in this message indicates the location of the database information used by the database manager. If space is unavailable in the path, the server can no longer function.

You must add space to the file system or make space available on the file system or disk.

Linux: Deleting installation rollback files

You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

About this task

To delete the files that are no longer needed, use either the installation graphical wizard or the command line in console mode.

- Linux: Deleting installation rollback files by using a graphical wizard
You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.
- Linux: Deleting installation rollback files by using the command line
You can delete certain installation files that were saved during the installation process by using the command line.

Linux: Deleting installation rollback files by using a graphical wizard

You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.

Procedure

1. Open IBM Installation Manager.

Linux In the directory where IBM Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command to start IBM Installation Manager:

```
./IBMIM
```

2. Click File > Preferences.
3. Select Files for Rollback.
4. Click Delete Saved Files and click OK.

Linux: Deleting installation rollback files by using the command line

You can delete certain installation files that were saved during the installation process by using the command line.

Procedure

1. In the directory where IBM® Installation Manager is installed, go to the following subdirectory:
 - o **Linux** eclipse/tools

For example:

- o **Linux** /opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command to start an IBM Installation Manager command line:
 - o **Linux** ./imcl -c
 3. Enter **P** to select Preferences.
 4. Enter **B** to select Files for Rollback.
 5. Enter **D** to Delete the Files for Rollback.
 6. Enter **A** to Apply Changes and Return to Preferences Menu.
 7. Enter **C** to leave the Preference Menu.
 8. Enter **X** to Exit Installation Manager.

Linux: Server naming best practices

Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect™ server.

Instance user ID

The instance user ID is used as the basis for other names related to the server instance. The instance user ID is also called the instance owner.

For example: tsminst1

The instance user ID is the user ID that must have ownership or read/write access authority to all directories that you create for the database and the recovery log. The standard way to run the server is under the instance user ID. That user ID must also have read/write access to the directories that are used for any FILE device classes.

Linux

Home directory for the instance user ID

The home directory can be created when creating the instance user ID, by using the option (-m) to create a home directory if it does not exist already. Depending on local settings, the home directory might have the form: /home/instance_user_ID

For example: /home/tsminst1

The home directory is primarily used to contain the profile for the user ID and for security settings.

Linux

Database instance name

The database instance name must be the same as the instance user ID under which you run the server instance.

For example: tsminst1

Linux

Instance directory

The instance directory is a directory that contains files specifically for a server instance (the server options file and other server-specific files). It can have any name that you want. For easier identification, use a name that ties the directory to the instance name.

You can create the instance directory as a subdirectory of the home directory for the instance user ID. For example:
`/home/instance_user_ID/instance_user_ID`

The following example places the instance directory in the home directory for user ID tsminst1: `/home/tsminst1/tsminst1`

You can also create the directory in another location, for example: `/tsmsserver/tsminst1`

The instance directory stores the following files for the server instance:

- The server options file, `dsmserv.opt`
- The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name
- Volume history file, if the `VOLUMEHISTORY` server option does not specify a fully qualified name
- Volumes for `DEVTYPE=FILE` storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)

Database name

The database name is always `TSMDB1`, for every server instance. This name cannot be changed.

Server name

The server name is an internal name for IBM Spectrum Protect, and is used for operations that involve communication among multiple IBM Spectrum Protect servers. Examples include server-to-server communication and library sharing.

Linux The server name is also used when you add the server to the Operations Center so that it can be managed using that interface. Use a unique name for each server. For easy identification in the Operations Center (or from a `QUERY SERVER` command), use a name that reflects the location or purpose of the server. Do not change the name of an IBM Spectrum Protect server after it is configured as a hub or spoke server.

If you use the wizard, the default name that is suggested is the host name of the system that you are using. You can use a different name that is meaningful in your environment. If you have more than one server on the system and you use the wizard, you can use the default name for only one of the servers. You must enter a unique name for each server.

Linux For example:

- `PAYROLL`
- `SALES`

Directories for database space and recovery log

The directories can be named according to local practices. For easier identification, consider using names that tie the directories to the server instance.

For example, for the archive log:

- `Linux /tsminst1_archlog`

Linux: Installation directories

Installation directories for the IBM Spectrum Protect™ server include the server, DB2®, device, language, and other directories. Each one contains several additional directories.

The (/opt/tivoli/tsm/server/bin) is the default directory that contains server code and licensing.

The DB2 product that is installed as part of the installation of the IBM Spectrum Protect server has the directory structure as documented in DB2 information sources. Protect these directories and files as you do the server directories. The default directory is /opt/tivoli/tsm/db2.

You can use US English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, and Russian.

Linux: Installing the server components

To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.

About this task

Using the IBM Spectrum Protect™ installation software, you can install the following components:

- server
Tip: The database (DB2®), the Global Security Kit (GSKit) and IBM® Java™ Runtime Environment (JRE) are automatically installed when you select the server component.
- server languages
- license
- devices
- IBM Spectrum Protect for SAN
- Operations Center

Linux Allow approximately 30 - 45 minutes to install a V8.1.5 server, using this guide.

- Linux: Obtaining the installation package
You can obtain the IBM Spectrum Protect installation package from an IBM download site such as Passport Advantage® or IBM Fix Central.
- Linux: Installing IBM Spectrum Protect by using the installation wizard
You can install the server by using the IBM Installation Manager graphical wizard.
- Linux: Installing IBM Spectrum Protect by using console mode
You can install IBM Spectrum Protect by using the command line in console mode.
- Linux: Installing IBM Spectrum Protect in silent mode
You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.
- Linux: Installing server language packages
Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Linux: Obtaining the installation package

You can obtain the IBM Spectrum Protect™ installation package from an IBM® download site such as Passport Advantage® or IBM Fix Central.

Linux

Before you begin

If you plan to download the files, set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly:

1. To query the maximum file size value, issue the following command:

```
ulimit -Hf
```

2. If the system user limit for maximum file size is not set to unlimited, change it to unlimited by following the instructions in the documentation for your operating system.

Procedure

1. Download the appropriate package file from one of the following websites.
 - o Download the server package from Passport Advantage or Fix Central.
 - o For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you downloaded the package from an IBM download site, complete the following steps:

Linux

- a. Verify that you have enough space to store the installation files when they are extracted from the product package. See the download document for the space requirements:
 - IBM Spectrum Protect technote 4042944
 - IBM Spectrum Protect Extended Edition technote 4042945
 - IBM Spectrum Protect for Data Retention technote 4042946
- b. Download the package file to the directory of your choice. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.
- c. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

- d. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file, for example:

Linux

```
8.1.x.000-IBM-SPSRV-Linuxx86_64.bin  
8.1.x.000-IBM-SPSRV-Linuxs390x.bin  
8.1.x.000-IBM-SPSRV-Linuxppc64le.bin
```

3. Select one of the following methods of installing IBM Spectrum Protect:
 - o Linux: Installing IBM Spectrum Protect by using the installation wizard
 - o Linux: Installing IBM Spectrum Protect by using console mode
 - o Linux: Installing IBM Spectrum Protect in silent mode
4. After you install IBM Spectrum Protect, and before you customize it for your use, go to the IBM Support Portal. Click Support and downloads and apply any applicable fixes.

Linux: Installing IBM Spectrum Protect by using the installation wizard

You can install the server by using the IBM® Installation Manager graphical wizard.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install IBM Spectrum Protect™ by using this method:

Option	Description
--------	-------------

Option	Description
Installing the software from a downloaded package:	a. Change to the directory where you downloaded the package. b. Start the installation wizard by issuing the following command: Linux <pre>./install.sh</pre>

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.
You can view installation log files by clicking File > View Log from the Installation Manager tool. To collect these log files, click Help > Export Data for Problem Analysis from the Installation Manager tool.
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- Linux** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Linux: Installing IBM Spectrum Protect by using console mode

You can install IBM Spectrum Protect™ by using the command line in console mode.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install IBM Spectrum Protect by using this method:

Option	Description
Installing the software from a downloaded package:	a. Change to the directory where you downloaded the package. b. Start the installation wizard in console mode by issuing the following command: Linux <pre>./install.sh -c</pre> Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the Summary panel, specify G to generate the responses.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - Linux** /var/ibm/InstallationManager/logs
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- Linux** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Linux: Installing IBM Spectrum Protect in silent mode

You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the input directory where the installation package is extracted:

install_response_sample.xml

Use this file to install the IBM Spectrum Protect™ components.

update_response_sample.xml

Use this file to upgrade the IBM Spectrum Protect components.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
2. If you install the server and Operations Center in silent mode, create a password for the Operations Center truststore in the response file.
If you are using the install_response_sample.xml file, add the password in the following line of the file, where *mypassword* represents the password:

```
<variable name='ssl.password' value='mypassword' />
```

For more information about this password, see Installation checklist

Tip: To upgrade the Operations Center, the truststore password is not required if you are using the update_response_sample.xml file.

3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:

o **Linux**

```
./install.sh -s -input response_file -acceptLicense
```

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - o **Linux** /var/ibm/InstallationManager/logs
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **Linux** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Linux

Linux: Installing server language packages

Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Before you begin

For instructions on installing storage agent language packages, see Language pack configuration for storage agents.

- **Linux:** Server language locales
Use either the default language package option or select another language package to display server messages and help.
- **Linux:** Configuring a language package
After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.
- **Linux:** Updating a language package
You can modify or update a language package by using the IBM® Installation Manager.

Linux: Server language locales

Use either the default language package option or select another language package to display server messages and help.

Linux This language package is automatically installed for the following default language option for IBM Spectrum Protect™ server messages and help:

- **Linux** LANGUAGE en_US

For languages or locales other than the default, install the language package that your installation requires.

You can use the languages that are shown:

Linux

Table 1. Server languages for Linux

LANGUAGE	LANGUAGE option value
Chinese, Simplified	zh_CN
	zh_CN.gb18030
	zh_CN.utf8
Chinese, Traditional	Big5 / Zh_TW
	zh_TW
	zh_TW.utf8
English, United States	en_US
	en_US.utf8
French	fr_FR
	fr_FR.utf8
German	de_DE
	de_DE.utf8
Italian	it_IT
	it_IT.utf8
Japanese	ja_JP
	ja_JP.utf8
Korean	ko_KR
	ko_KR.utf8
Portuguese, Brazilian	pt_BR
	pt_BR.utf8
Russian	ru_RU
	ru_RU.utf8
Spanish	es_ES
	es_ES.utf8

Linux Restriction: For Operations Center users, some characters might not be displayed properly if the web browser does not use the same language as the server. If this problem occurs, set the browser to use the same language as the server.

Linux: Configuring a language package

After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.

About this task

Linux To set support for a certain locale, complete one of the following tasks:

- Set the LANGUAGE option in the server options file to the name of the locale that you want to use. For example:
 - **Linux** To use the `it_IT` locale, set the LANGUAGE option to `it_IT`. See [Linux: Server language locales](#).
- **Linux** If you are starting the server in the foreground, set the `LC_ALL` environment variable to match the value that is set in the server options file. For example, to set the environment variable for Italian, enter the following value:

```
export LC_ALL=it_IT
```

If the locale is successfully initialized, it formats the date, time, and number for the server. If the locale is not successfully initialized, the server uses the US English message files and the date, time, and number format.

Linux: Updating a language package

You can modify or update a language package by using the IBM® Installation Manager.

About this task

You can install another language package within the same IBM Spectrum Protect™ instance.

- Use the Modify function of IBM Installation Manager to install another language package.
- Use the Update function of IBM Installation Manager to update to newer versions of the language packages.

Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.

Linux: Taking the first steps after you install IBM Spectrum Protect

After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect™ instance.

About this task

1. **Linux** Update the kernel parameter values.
 - **Linux** See [Linux: Tuning kernel parameters for Linux systems](#).
 2. Create the directories and user ID for the server instance. See [Linux: Creating the user ID and directories for the server instance](#).
 3. Configure a server instance. Select one of the following options:
 - Use the configuration wizard, the preferred method. See [Linux: Configuring IBM Spectrum Protect by using the configuration wizard](#).
 - Manually configure the new instance. See [Linux: Configuring the server instance manually](#). Complete the following steps during a manual configuration.
 - a. Set up your directories and create the IBM Spectrum Protect instance. See [Linux: Creating the server instance](#).
 - b. Create a new server options file by copying the sample file to set up communications between the server and clients. See **Linux** [Linux: Configuring server and client communications](#).
 - c. Issue the `DSMSERV FORMAT` command to format the database. See [Linux: Formatting the database and log](#).
 - d. Configure your system for database backup. See [Linux: Preparing the database manager for database backup](#).
 4. Configure options to control when database reorganization runs. See [Linux: Configuring server options for server database maintenance](#).
 5. Start the server instance if it is not already started.
 - **Linux** See [Linux: Starting the server instance](#).
 6. Register your license. See [Linux: Registering licenses](#).
 7. Prepare your system for database backups. See [Linux: Preparing the server for database backup operations](#).
 8. Monitor the server. See [Linux: Monitoring the server](#).
- **Linux** [Linux: Tuning kernel parameters for Linux systems](#)
For IBM Spectrum Protect and DB2 to install and operate correctly on Linux, you must update the kernel configuration parameters.

- **Linux:** Creating the user ID and directories for the server instance
Create the user ID for the IBM Spectrum Protect server instance and create the directories that the server instance needs for database and recovery logs.
- **Linux:** Configuring the IBM Spectrum Protect server
After you have installed the server and prepared for the configuration, configure the server instance.
- **Linux:** Configuring server options for server database maintenance
To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.
- **Linux:** Starting the server instance
You can start the server by using the instance user ID, which is the preferred method, or the root user ID.
- **Linux:** Stopping the server
You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.
- **Linux:** Registering licenses
Immediately register any IBM Spectrum Protect licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.
- **Linux:** Preparing the server for database backup operations
To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.
- **Linux:** Running multiple server instances on a single system
You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.
- **Linux:** Monitoring the server
When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Linux

Linux: Tuning kernel parameters for Linux systems

For IBM Spectrum Protect™ and DB2® to install and operate correctly on Linux, you must update the kernel configuration parameters.

About this task

If you do not update these parameters, the installation of DB2 and IBM Spectrum Protect might fail. Even if installation is successful, operational problems might occur if you do not set parameter values.

- **Linux:** Updating kernel parameters on Linux
DB2 automatically increases interprocess communication (IPC) kernel parameter values to the preferred settings.
- **Linux:** Suggested values for kernel parameters on Linux
Ensure that the values for kernel parameters are sufficient to prevent operational problems from occurring when you run the IBM Spectrum Protect server.

Linux

Linux: Updating kernel parameters on Linux

DB2® automatically increases interprocess communication (IPC) kernel parameter values to the preferred settings.

About this task

To update the kernel parameters on Linux servers, complete the following steps:

Procedure

1. Issue the `ipcs -l` command to list the parameter values.
2. Analyze the results to determine whether any changes are required for your system. If changes are required, you can set the parameter in the `/etc/sysctl.conf` file. The parameter value is applied when the system starts.

What to do next

For Red Hat Enterprise Linux 6 (RHEL6), you must set the `kernel.shmmax` parameter in the `/etc/sysctl.conf` file before automatically starting the IBM Spectrum Protect™ server on system startup.

For details about the DB2 database for Linux, see the DB2 product information.

Linux

Linux: Suggested values for kernel parameters on Linux

Ensure that the values for kernel parameters are sufficient to prevent operational problems from occurring when you run the IBM Spectrum Protect™ server.

About this task

The following table contains the suggested kernel parameter settings to run both IBM Spectrum Protect and DB2®.

Parameter	Description	Preferred value
<code>kernel.randomize_va_space</code>	The <code>kernel.randomize_va_space</code> parameter configures the kernel's use of memory ASLR. When you set the value to 0, <code>kernel.randomize_va_space=0</code> , it disables ASLR. DB2 data servers rely on fixed addresses for certain shared memory objects, and the ASLR can cause errors for some activities. To learn more details about the Linux ASLR and DB2, see the technote at: http://www.ibm.com/support/docview.wss?uid=swg21365583 .	0
<code>vm.swappiness</code>	The <code>vm.swappiness</code> parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the DB2 product information.	0
<code>vm.overcommit_memory</code>	The <code>vm.overcommit_memory</code> parameter influences how much virtual memory the kernel can permit be allocated. For more information about kernel parameters, see the DB2 product information.	0

Linux: Creating the user ID and directories for the server instance

Create the user ID for the IBM Spectrum Protect™ server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See *Linux: Worksheets for planning details for the server*.

Procedure

1. Create the user ID that will own the server instance. You use this user ID when you create the server instance in a later step.

Linux

Linux Create a user ID and group that will be the owner of the server instance.

- a. The following commands can be run from an administrative user ID that will set up the user and group. Create the user ID and group in the home directory of the user.

Restriction: In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID and group name must comply with the following rules:

- The length must be 8 characters or less.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

For example, create user ID `tsminst1` in group `tsmsrvrs`. The following examples show how to create this user ID and group using operating system commands.

Linux

```
groupadd tsmsrvrs -g 1111
useradd -d /home/tsminst1 -u 2222 -g 1111 -s /bin/bash tsminst1
passwd tsminst1
```

Restriction: DB2® does not support direct operating system user authentication through LDAP.

- b. Log off, then log in to your system. Change to the user account that you just created. Use an interactive login program, such as telnet, so that you are prompted for the password and can change it if necessary.

2. Create directories that the server requires.

Linux

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	<code>mkdir /tsminst1</code>	
The database directories	<code>mkdir /tsmdb001</code> <code>mkdir /tsmdb002</code> <code>mkdir /tsmdb003</code> <code>mkdir /tsmdb004</code>	
Active log directory	<code>mkdir /tsmlog</code>	
Archive log directory	<code>mkdir /tsmarchlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir /tsmlogmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir /tsmarchlogfailover</code>	

When a server is initially created by using the DSMSEV FORMAT utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

Linux: Configuring the IBM Spectrum Protect server

After you have installed the server and prepared for the configuration, configure the server instance.

About this task

Configure an IBM Spectrum Protect™ server instance by selecting one of the following options:

- **Linux: Configuring IBM Spectrum Protect by using the configuration wizard**
The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect server program.
- **Linux: Configuring the server instance manually**
After installing IBM Spectrum Protect Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the

configuration wizard.

Linux: Configuring IBM Spectrum Protect by using the configuration wizard

The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect™ server program.

Before you begin

Before you use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing IBM Spectrum Protect, creating the database and log directories, and creating the directories and user ID for the server instance.

Procedure

1. Ensure that the following requirements are met:

Linux

- The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/` directory. Also, ensure that the SSH daemon service has access rights for connecting to the system by using the `localhost` value.
- You must be able to log in to the system with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- A backup copy of the following files must be saved to a safe and secure location:
 - Master encryption key files (`dsmkeydb.*`)
 - Server certificate and private key files (`cert.*`)

2. Start the local version of the wizard:

- **Linux** Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be run only by using the root user ID.

Follow the instructions to complete the configuration. The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Linux: Configuring the server instance manually

After installing IBM Spectrum Protect™ Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

- Linux: Creating the server instance
Create an IBM Spectrum Protect instance by issuing the `db2icrt` command.
- **Linux** Linux: Configuring server and client communications
A default sample server options file, `dsm serv.opt.smp`, is created during IBM Spectrum Protect installation in the `/opt/tivoli/tsm/server/bin` directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.
- Linux: Formatting the database and log
Use the `DSMSERV FORMAT` utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.
- Linux: Preparing the database manager for database backup
To back up the data in the database to IBM Spectrum Protect, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

Linux: Creating the server instance

Create an IBM Spectrum Protect™ instance by issuing the `db2icrt` command.

About this task

You can have one or more server instances on one workstation.

Linux Important: Before you run the db2icrt command, verify the following items:

- The home directory for the user (/home/tsminst1) exists. If there is no home directory, you must create it. The instance directory stores the following files that are generated by the IBM Spectrum Protect server:
 - The server options file, dsmserv.opt
 - The server key database file, cert.kdb, and the .arm files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
 - Device configuration file, if the DEVCONFIG server option does not specify a fully qualified name
 - Volume history file, if the VOLUMEHISTORY server option does not specify a fully qualified name
 - Volumes for DEVTYPE=FILE storage pools, if the directory for the device class is not fully specified, or not fully qualified
 - User exits
 - Trace output (if not fully qualified)
- A backup copy of the following files must be saved to a safe and secure location:
 - Master encryption key files (dsmkeydb.*)
 - Server certificate and private key files (cert.*)
- The root user and instance-user ID must have write permission to the shell configuration file. A shell configuration file (for example, .profile) exists in the home directory. For more information, see the DB2® product information. Search for Linux and UNIX environment variable settings.

Linux

1. Log in using the root user ID and create an IBM Spectrum Protect instance. The name of the instance must be the same name as the user that owns the instance. Use the db2icrt command and enter the command on one line: **Linux**

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
instance_name instance_name
```

For example, if your user ID for this instance is tsminst1, use the following command to create the instance. Enter the command on one line. **Linux**

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
tsminst1 tsminst1
```

Remember: From this point on, use this new user ID when you configure your IBM Spectrum Protect server. Log out of the root user ID and log in under the new instance-user ID.

2. Change the default directory for the database to be the same as the instance directory for the server. If you have multiple servers, log in under the instance ID for each server. Issue this command:

```
db2 update dbm cfg using dftdbpath instance_directory
```

For example, where instance_directory is the instance user ID:

```
db2 update dbm cfg using dftdbpath /tsminst1
```

3. Modify the library path to include libraries that are required for server operations.

Tip: In the following examples, here are the directories:

- *server_bin_directory* is a subdirectory of the server installation directory. For example, /opt/tivoli/tsm/server/bin.
- *instance_users_home_directory* is the home directory of the instance user. For example, /home/tsminst1.

- You must update one of the following files to set the library path when DB2 or the server are started. Update per the shell that the instance user is configured to use.

Bash or Korn shell:

```
instance_users_home_directory/sqlllib/userprofile
```

C shell:

```
instance_users_home_directory/sqlllib/usercshrc
```

- Update per the shell that the instance user is configured to use.

Bash or Korn shell:

Add the following entry to the `instance_users_home_directory/sqlllib/userprofile` file, on one line: **Linux**

```
export LD_LIBRARY_PATH=server_bin_directory/  
dbbkapi:/usr/local/ibm/gsk8_64/lib64:/opt/ibm/lib:/opt/  
ibmlib64:$LD_LIBRARY_PATH
```

C shell:

Add the following entry to the `instance_users_home_directory/sqlllib/usercshrc` file, on one line: **Linux**

```
setenv LD_LIBRARY_PATH server_bin_directory/dbbkapi:/  
usr/local/ibm/gsk8_64/lib64:/  
opt/ibm/lib:/opt/ibm/lib64:/usr/lib64:$LD_LIBRARY_PATH
```

Remember: The following entries must be in the library path, preceding any other entries in the library path:

- `server_bin_directory/dbbkapi`
- `/usr/local/ibm/gsk8_64/lib64`

4. Create a new server options file. See **Linux: Configuring server and client communications**.

Linux

Linux: Configuring server and client communications

A default sample server options file, `dsmserv.opt.smp`, is created during IBM Spectrum Protect™ installation in the `/opt/tivoli/tsm/server/bin` directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.

About this task

Ensure that you have a server instance directory, for example `/tsminst1`, and copy the sample file to this directory. Name the new file `dsmserv.opt` and edit the options. Complete this set-up before you initialize the server database. Each sample or default entry in the sample options file is a comment, a line beginning with an asterisk (*). Options are not case-sensitive and one or more blank spaces are allowed between keywords and values.

When editing the options file, follow these guidelines:

- Remove the asterisk at the beginning of the line to activate an option.
- Begin entering the options in any column.
- Enter only one option per line, and the option must be on only one line.
- If you make multiple entries for a keyword, the IBM Spectrum Protect server uses the last entry.

If you change the server options file, you must restart the server for the changes to take effect.

You can specify one or more of the following communication methods:

- TCP/IP Version 4 or Version 6
- Shared memory
- Secure Sockets Layer (SSL)
Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the IBM Spectrum Protect server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.
- **Linux** **Linux: Setting TCP/IP options**
Select from a range of TCP/IP options for the IBM Spectrum Protect server or retain the default.
- **Linux** **Linux: Setting shared memory options**
You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.
- **Linux** **Linux: Setting Secure Sockets Layer options**
You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Linux: Setting TCP/IP options

Select from a range of TCP/IP options for the IBM Spectrum Protect™ server or retain the default.

About this task

The following is an example of a list of TCP/IP options that you can use to set up your system.

```
commmethod      tcpip
tcpport         1500
tcpwindowsize   0
tcpnodelay      yes
```

Tip: You can use TCP/IP Version 4, Version 6, or both.

TCPPORT

The server port address for TCP/IP and SSL communication. The default value is 1500.

Linux TCPWINDOWSIZE

Specifies the size of the TCP/IP buffer that is used when sending or receiving data. The window size that is used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but can improve performance.

You can specify an integer from 0 to 2048. To use the default window size for the operating system, specify 0.

TCPNODELAY

Specifies whether or not the server sends small messages or lets TCP/IP buffer the messages. Sending small messages can improve throughput but increases the number of packets sent over the network. Specify YES to send small messages or NO to let TCP/IP buffer them. The default is YES.

TCPADMINPORT

Specifies the port number on which the server TCP/IP communication driver is to wait for TCP/IP or SSL-enabled communication requests other than client sessions. The default is the value of TCPPORT.

SSLTCPSPORT

(SSL-only) Specifies the Secure Sockets Layer (SSL) port number on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line backup-archive client and the command-line administrative client.

SSLTCPADMINPORT

(SSL-only) Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client.

Linux: Setting shared memory options

You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.

About this task

The following example shows a shared memory setting:

```
commmethod      sharedmem
shmport         1510
```

In this example, SHMPORT specifies the TCP/IP port address of a server when using shared memory. Use the SHMPORT option to specify a different TCP/IP port. The default port address is 1510.

COMMETHOD can be used multiple times in the IBM Spectrum Protect™ server options file, with a different value each time. For example, the following example is possible:

```
commmethod tcpip
commmethod sharedmem
```

Linux You might receive the following message from the server when using shared memory:

```
ANR9999D shmcomm.c(1598): ThreadId<39>
Error from msgget (2), errno = 28
```

The message means that a message queue must be created but the system limit for the maximum number of message queues (MSGMNI) would be exceeded.

Linux To find out the maximum number of message queues (MSGMNI) on your system, issue the following command:

```
cat /proc/sys/kernel/msgmni
```

To increase the MSGMNI value on your system, issue the following command:

```
sysctl -w kernel.msgmni=n
```


where **n** is the maximum number of message queues that you want the system to allow.

Linux: Setting Secure Sockets Layer options

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Before you begin

SSL is the standard technology for creating encrypted sessions between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communication paths. With SSL, the identity of the server is verified through the use of digital certificates.

To ensure better system performance, use SSL only for sessions when it is needed. Consider adding additional processor resources on the IBM Spectrum Protect™ server to manage the increased requirements.

Linux: Formatting the database and log

Use the DSMSEV FORMAT utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

After you set up server communications, you are ready to initialize the database. Ensure that you log in by using the instance user ID. Do not place the directories on file systems that might run out of space. If certain directories (for example, the archive log) become unavailable or full, the server stops.

Setting the exit list handler

Set the DB2NOEXITLIST registry variable to ON for each server instance. Log on to the system as the server instance owner and issue this command:

```
db2set -i server_instance_name DB2NOEXITLIST=ON
```

For example: **Linux**

```
db2set -i tsminst1 DB2NOEXITLIST=ON
```

Initializing a server instance

Use the DSMSEV FORMAT utility to initialize a server instance. For example, if the server instance directory is */tsminst1*, issue the following commands: **Linux**

```
cd /tsminst1
dmserv format dbdir=/tsmdb001 activelogsiz=32768
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

Linux Tip: If DB2® does not start after you issue the DSMSEV FORMAT command, you might need to disable the file system mount option NOSUID. If this option is set on the file system that contains the DB2 instance owner directory, or on any file system that contains the DB2 database, active logs, archive logs, failover logs, or mirrored logs, the option must be disabled to start the system.

After you disable the NOSUID option, remount the file system and then start DB2 by issuing the following command:

```
db2start
```

Related information:

[DSMSEV FORMAT \(Format the database and log\)](#)

Linux: Preparing the database manager for database backup

To back up the data in the database to IBM Spectrum Protect™, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

About this task

Linux Starting with IBM Spectrum Protect V7.1, it is no longer necessary to set the API password during a manual configuration of the server. If you set the API password during the manual configuration process, attempts to back up the database might fail.

If you use the configuration wizard to create an IBM Spectrum Protect server instance, you do not have to complete these steps. If you are configuring an instance manually, complete the following steps before you issue either the BACKUP DB or the RESTORE DB commands.

Attention: If the database is unusable, the entire IBM Spectrum Protect server is unavailable. If a database is lost and cannot be recovered, it might be difficult or impossible to recover data that is managed by that server. Therefore, it is critically important to back up the database.

Linux In the following commands, replace the example values with your actual values. The examples use `tsminst1` for the server instance user ID, `/tsminst1` for the server instance directory, and `/home/tsminst1` as the server instance users home directory.

1. Set the IBM Spectrum Protect API environment-variable configuration for the database instance:

- a. Log in by using the `tsminst1` user ID.
- b. When user `tsminst1` is logged in, ensure that the DB2® environment is properly initialized. The DB2 environment is initialized by running the `/home/tsminst1/sqllib/db2profile` script, which normally runs automatically from the profile of the user ID. Ensure the `.profile` file exists in the instance users home directory, for example, `/home/tsminst1/.profile`. If `.profile` does not run the `db2profile` script, add the following lines:

```
if [ -f /home/tsminst1/sqllib/db2profile ]; then
    . /home/tsminst1/sqllib/db2profile
fi
```

- c. In the `instance_directory/sqllib/userprofile` file, add the following lines:

```
DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
DSMI_DIR=server_bin_directory/dbbkapi
DSMI_LOG=server_instance_directory
export DSMI_CONFIG DSMI_DIR DSMI_LOG
```

where:

- *instance_directory* is the home directory of the server instance user.
- *server_instance_directory* is the server instance directory.
- *server_bin_directory* is the server bin directory. The default location is `/opt/tivoli/tsm/server/bin`.

In the `instance_directory/sqllib/usercshrc` file, add the following lines:

```
setenv DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
setenv DSMI_DIR=server_bin_directory/dbbkapi
setenv DSMI_LOG=server_instance_directory
```

2. Log off and log in again as `tsminst1`, or issue this command:

```
. ~/.profile
```

Tip: Ensure that you enter a space after the initial dot (.) character.

3. Create a file that is named `tsmdbmgr.opt` in the *server_instance* directory, which is in the `/tsminst1` directory in this example, and add the following line:

```
SERVERNAME TSMDBMGR_TSMINST1
```

Remember: The value for `SERVERNAME` must be consistent in the `tsmdbmgr.opt` and `dsm.sys` files.

4. As root user, add the following lines to the IBM Spectrum Protect API `dsm.sys` configuration file. By default, the `dsm.sys` configuration file is in the following default location:

- o `server_bin_directory/dbbkapi/dsm.sys`

```
servername TSMDBMGR_TSMINST1
commethod tcpip
tcpserveraddr localhost
tcpport 1500
errorlogname /tsminst1/tsmdbmgr.log
nodename $$TSMDBMGR_$$
```

where

- o *servername* matches the *servername* value in the `tsmdbmgr.opt` file.
 - o *commethod* specifies the client API that is used to contact the server for database backup. This value can be `tcpip` or `sharedmem`. For more information about shared memory, see step 5.
 - o *tcpserveraddr* specifies the server address that the client API uses to contact the server for database backup. To ensure that the database can be backed up, this value must be `localhost`.
 - o *tcpport* specifies the port number that the client API uses to contact the server for database backup. Ensure that you enter the same `tcpport` value that is specified in the `dsmserv.opt` server options file.
 - o *errorlogname* specifies the error log where the client API logs errors that are encountered during a database backup. This log is typically in the server instance directory. However, this log can be placed in any location where the instance user ID has write-permission.
 - o *nodename* specifies the node name that the client API uses to connect to the server during a database backup. To ensure that the database can be backed up, this value must be `$$_TSMDBMGR_$$`.
- Linux** Attention: Do not add the `PASSWORDACCESS generate` option to the `dsm.sys` configuration file. This option can cause the database backup to fail.
5. Optional: Configure the server to back up the database by using shared memory. In this way, you might be able to reduce the processor load and improve throughput. Complete the following steps:
- a. Review the `dsmserv.opt` file. If the following lines are not in the file, add them:

```
commethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

- b. In the `dsm.sys` configuration file, locate the following lines:

```
commethod tcpip
tcpserveraddr localhost
tcpport port_number
```

Replace the specified lines with the following lines:

```
commethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

Linux: Configuring server options for server database maintenance

To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.

About this task

Table and index reorganization requires significant processor resources, active log space, and archive log space. Because database backup takes precedence over reorganization, select the time and duration for reorganization to ensure that the processes do not overlap and reorganization can complete.

Linux You can optimize index and table reorganization for the server database. In this way, you can help to avoid unexpected database growth and performance issues. For instructions, see technote 1683633.

If you update these server options while the server is running, you must stop and restart the server before the updated values take effect.

Procedure

1. Modify the server options.

Linux Edit the server options file, `dsmserv.opt`, in the server instance directory. Follow these guidelines when you edit the server options file:

- o To enable an option, remove the asterisk at the beginning of the line.
- o Enter an option on any line.
- o Enter only one option per line. The entire option with its value must be on one line.
- o If you have multiple entries for an option in the file, the server uses the last entry.

To view available server options, see the sample file, `dsmserv.opt.smp`, in the `/opt/tivoli/tsm/server/bin` directory.

2. If you plan to use data deduplication, enable the ALLOWREORGINDEX server option. Add the following option and value to the server options file:

```
allowreorgindex yes
```

3. Set the REORGBEGINTIME and REORGDURATION server options to control when reorganization starts and how long it runs. Select a time and duration so that reorganization runs when you expect that the server is least busy. These server options control both table and index reorganization processes.
 - a. Set the time for reorganization to start by using the REORGBEGINTIME server option. Specify the time by using the 24-hour system. For example, to set the start time for reorganization as 8:30 p.m., specify the following option and value in the server options file:

```
reorgbegintime 20:30
```

- b. Set the interval during which the server can start reorganization. For example, to specify that the server can start reorganization for four hours after the time set by the REORGBEGINTIME server option, specify the following option and value in the server options file:

```
reorgduration 4
```

4. If the server was running while you updated the server options file, stop and restart the server.

Linux

Linux: Starting the server instance

You can start the server by using the instance user ID, which is the preferred method, or the root user ID.

Before you begin

Ensure that you set access permissions and user limits correctly.

Linux For instructions, see [Verifying access rights and user limits](#).

About this task

When you start the server by using the instance user ID, you simplify the setup process and avoid potential issues. However, in some cases, it might be necessary to start the server with the root user ID. For example, you might want to use the root user ID to ensure that the server can access specific devices. You can set up the server to start automatically by using either the instance user ID or the root user ID.

Linux If you must complete maintenance or reconfiguration tasks, start the server in maintenance mode.

Procedure

To start the server, take one of the following actions:

- Start the server by using the instance user ID.

Linux For instructions, see [Starting the server from the instance user ID](#).

- Start the server by using the root user ID.

For instructions about authorizing root user IDs to start the server, see [Authorizing root user IDs to start the server \(V7.1.1\)](#). For instructions about starting the server by using the root user ID, see [Starting the server from the root user ID \(V7.1.1\)](#).

- **Linux** Start the server automatically.

Linux For instructions, see [Linux: Automatically starting servers on Linux systems](#).

- **Linux** Start the server in maintenance mode.

For instructions, see [Linux: Starting the server in maintenance mode](#).

Linux

Linux: Verifying access rights and user limits

Before you start the server, verify access rights and user limits.

About this task

If you do not verify user limits, also known as *ulimits*, you might experience server instability or a failure of the server to respond. You must also verify the system-wide limit for the maximum number of open files. The system-wide limit must be greater than or equal to the user limit.

Procedure

1. Verify that the server instance user ID has permissions to start the server.
2. For the server instance that you plan to start, ensure that you have authority to read and write files in the server instance directory. Verify that the `dsmserv.opt` file exists in the server instance directory, and that the file includes parameters for the server instance.
3. If the server is attached to a tape drive, medium changer, or removable media device, and you plan to start the server by using the instance user ID, grant read/write access to the instance user ID for these devices. To set permissions, take one of the following actions:

- o If the system is dedicated to IBM Spectrum Protect™ and only the IBM Spectrum Protect administrator has access, make the device special file world-writable. On the operating system command line, issue the following command:

```
chmod +w /dev/rmtX
```

- o If the system has multiple users, you can restrict access by making the IBM Spectrum Protect instance user ID the owner of the special device files. On the operating system command line, issue the following command:

```
chmod u+w /dev/rmtX
```

- o If multiple user instances are running on the same system, change the group name, for example TAPEUSERS, and add each IBM Spectrum Protect instance user ID to that group. Then, change the ownership of the device special files to belong to the group TAPEUSERS, and make them group-writable. On the operating system command line, issue the following command:

```
chmod g+w /dev/rmtX
```

4. **Linux** If you are using the IBM Spectrum Protect device driver and the `autoconf` utility, use the `-a` option to grant read/write access to the instance user ID.

5. **Linux** To prevent server failures during interaction with DB2®, tune the kernel parameters.

Linux For instructions about tuning kernel parameters, see [Linux: Tuning kernel parameters for Linux systems](#).

6. Verify the following user limits based on the guidelines in the table.

Table 1. User limit (ulimit) values

User limit type	Preferred value	Command to query value
Maximum size of core files created	Unlimited	<code>ulimit -Hc</code>
Maximum size of a data segment for a process	Unlimited	<code>ulimit -Hd</code>
Maximum file size	Unlimited	<code>ulimit -Hf</code>
Maximum number of open files	65536	<code>ulimit -Hn</code>
Maximum amount of processor time in seconds	Unlimited	<code>ulimit -Ht</code>

To modify user limits, follow the instructions in the documentation for your operating system.

Tip: If you plan to start the server automatically by using a script, you can set the user limits in the script.

7. Ensure that the user limit of maximum user processes (the `nproc` setting) is set to the minimum suggested value of 16384.
 - a. To verify the current user limit, issue the `ulimit -Hu` command by using the instance user ID. For example:

```
[user@Machine ~]$ ulimit -Hu
16384
```

- b. If the limit of maximum user processes is not set to 16384, set the value to 16384.

Linux Add the following line to the `/etc/security/limits.conf` file:

```
instance_user_id          -      nproc           16384
```

where `instance_user_id` specifies the server instance user ID.

Linux If the server is installed on the Red Hat Enterprise Linux 6 operating system, set the user limit by editing the `/etc/security/limits.d/90-nproc.conf` file in the `/etc/security/limits.d` directory. This file overrides the settings in the `/etc/security/limits.conf` file.

Tip: The default value for the user limit of maximum user processes has changed on some distributions and versions of the Linux operating system. The default value is 1024. If you do not change the value to the minimum suggested value of 16384, the server might fail or hang.

Linux

Linux: Starting the server from the instance user ID

To start the server from the instance user ID, log in with the instance user ID and issue the appropriate command from the server instance directory.

Before you begin

Ensure that access rights and user limits are set correctly. For instructions, see [Linux: Verifying access rights and user limits](#).

Procedure

1. Log in to the system where IBM Spectrum Protect™ is installed by using the instance user ID for the server.
2. If you do not have a user profile that runs the `db2profile` script, issue the following command:

```
. /home/tsminst1/sqlllib/db2profile
```

Tip: For instructions about updating the user ID login script to run the `db2profile` script automatically, see the DB2® documentation.

3. Start the server by issuing the following command on one line from the server instance directory:

Linux

```
usr/bin/dmserv
```

Tip: The command runs in the foreground so that you can set an administrator ID and connect to the server instance.

Linux For example, if the name of the server instance is `tsminst1` and the server instance directory is `/tsminst1`, you can start the instance by issuing the following commands:

```
cd /tsminst1
. ~/sqlllib/db2profile
/usr/bin/dmserv
```

Linux

Linux: Automatically starting servers on Linux systems

To automatically start a server on a Linux operating system, use the `dmserv.rc` script.

Before you begin

Ensure that kernel parameters are set correctly. For instructions, see [Tuning kernel parameters for Linux systems](#).

Ensure that the server instance runs under the instance owner user ID.

Ensure that access rights and user limits are set correctly. For instructions, see [Verifying access rights and user limits](#).

About this task

The `dmserv.rc` script is in the server installation directory, for example, `/opt/tivoli/tsm/server/bin`.

The `dsmserv.rc` script can be used either to start the server manually or to start the server automatically by adding entries to the `/etc/rc.d/init.d` directory. The script works with Linux utilities such as `CHKCONFIG` and `SERVICE`.

Procedure

For each server instance that you want to automatically start, complete the following steps:

1. Place a copy of the `dsmserv.rc` script in the `/init.d` directory, for example, `/etc/rc.d/init.d`.

Ensure that you change only the copy of the script. Do not change the original script.

2. Rename the script copy so that it matches the name of the server instance owner, for example, `tsminst1`.

The script was created under the assumption that the server instance directory is `home_directory/tsminst1`, for example: `/home/tsminst1/tsminst1`.

3. If the server instance directory is not `home_directory/tsminst1`, locate the following line in the script copy:

```
instance_dir="${instance_home}/tsminst1"
```

Change the line so that it points to your server instance directory, for example:

```
instance_dir="/tsminst1"
```

4. In the script copy, locate the following line:

```
# pidfile: /var/run/dsmserv_instancename.pid
```

Change the instance name value to the name of the server instance owner. For example, if the server instance owner is `tsminst1`, update the line as shown:

```
# pidfile: /var/run/dsmserv_tsminst1.pid
```

5. Configure the run level in which the server automatically starts. By using tools such as the `CHKCONFIG` utility, specify a value that corresponds to a multiuser mode, with networking turned on. Typically, the run level to use is 3 or 5, depending on the operating system and its configuration. For more information about multiuser mode and run levels, see the documentation for your operating system.
6. To start or stop the server, issue one of the following commands:

- o To start the server:

```
service tsminst1 start
```

- o To stop the server:

```
service tsminst1 stop
```

Example

This example uses the following values:

- The instance owner is `tsminst1`.
- The server instance directory is `/home/tsminst1/tsminst1`.
- The `dsmserv.rc` script copy is named `tsminst1`.
- The `CHKCONFIG` utility is used to configure the script to start at run levels 3, 4, and 5.

```
cp /opt/tivoli/tsm/server/bin/dsmserv.rc /etc/rc.d/init.d/tsminst1
sed -i 's/dsmserv_instancename.pid/dsmserv_tsminst1.pid/' /etc/rc.d/init.d/tsminst1
chkconfig --list tsminst1
service tsminst1 supports chkconfig, but is not referenced in
any runlevel (run 'chkconfig --add tsminst1')
chkconfig --add tsminst1
chkconfig --list tsminst1
tsminst1 0:off 1:off 2:off 3:off 4:off 5:off 6:off
chkconfig --level 345 tsminst1 on
chkconfig --list tsminst1
tsminst1 0:off 1:off 2:off 3:on 4:on 5:on 6:off
```

Related information:

[Server startup script: dsmserv.rc](#)

Linux: Starting the server in maintenance mode

You can start the server in maintenance mode to avoid disruptions during maintenance and reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSErv utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode.

Operations that were disabled during maintenance mode are reenabled.

Linux: Stopping the server

You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.

About this task

To stop the server, issue the following command from the IBM Spectrum Protect™ command line:

```
halt
```

Linux If you cannot connect to the server with an administrative client and you want to stop the server, you must cancel the process by using the kill command with the process ID number (pid). The pid is displayed at initialization.

Important: Before you issue the kill command, ensure that you know the correct process ID for the IBM Spectrum Protect server. The `dsmserv.v6lock` file, in the directory from which the server is running, can be used to identify the process ID of the process to kill. To display the file, enter:

```
cat /instance_dir/dsmserv.v6lock
```


Linux Issue the following command to stop the server:

```
kill -23 dmserv_pid
```

where *dmserv_pid* is the process ID number.

Linux: Registering licenses

Immediately register any IBM Spectrum Protect™ licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.

About this task

Use the REGISTER LICENSE command for this task. See REGISTER LICENSE for more details.

Example: Register a license

Register the base IBM Spectrum Protect license.

```
register license file=tsmbasic.lic
```

Linux: Preparing the server for database backup operations

To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.

Procedure

1. Ensure that the IBM Spectrum Protect™ configuration is complete. If you did not use the configuration wizard (dsmicfgx) to configure the server, ensure that you completed the steps to manually configure the server for database backups.
2. Select the device class to be used for database backups, protect the master encryption key, and set a password. All of these actions are completed by issuing the SET DBRECOVERY command from the administrative command line:

```
set dbrecovery device_class_name protectkeys=yes password=password_name
```

where *device_class_name* specifies the device class to be used for database backup operations, and *password_name* specifies the password.

You must specify a device class name or the backup fails. By specifying PROTECTKEYS=YES, you ensure that the master encryption key is backed up during database backup operations.

Important: Create a strong password that is at least 8 characters long. Ensure that you remember this password. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database.

Example

To specify that database backups include a copy of the master encryption key for the server, run the following command:

```
set dbrecovery dbback protectkeys=yes password=protect8991
```

Linux: Running multiple server instances on a single system

You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

Multiply the memory and other system requirements for one server by the number of instances planned for the system.

Linux The set of files for one instance of the server is stored separately from the files used by another server instance on the same system. Use the steps in Linux: Creating the server instance for each new instance, including creation of the new instance user.

To manage the system memory that is used by each server, use the DBMEMPERCENT server option to limit the percentage of system memory. If all servers are equally important, use the same value for each server. If one server is a production server and

other servers are test servers, set the value for the production server to a higher value than the test servers.

You can upgrade directly from V7.1 to V8.1. See the upgrade section (Upgrading to V8.1) for more details. When you upgrade and have multiple servers on your system, you must run the installation wizard only once. The installation wizard collects the database and variables information for all of your original server instances.

If you upgrade from IBM Spectrum Protect V6.3 to V8.1.5 and have multiple servers on your system, all instances that exist in DB2® V9.7 are dropped and recreated in DB2 V11.1. The wizard issues the `db2 upgrade db dbname` command for each database. The database environment variables for each instance on your system are also reconfigured during the upgrade process.

Related tasks:

[Running multiple server instances on a single system \(V7.1.1\)](#)

Linux: Monitoring the server

When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Procedure

1. Monitor the active log to ensure that the size is correct for the workload that is handled by the server instance.

When the server workload reaches its typical expected level, the space that is used by the active log is 80% - 90% of the space that is available to the active log directory. At that point, you might need to increase the amount of space. Whether you must increase the space depends on the types of transactions in the server workload. Transaction characteristics affect how the active log space is used.

The following transaction characteristics can affect the space usage in the active log:

- o The number and size of files in backup operations
 - Clients such as file servers that back up large numbers of small files can cause large numbers of transactions that are completed quickly. The transactions might use a large amount of space in the active log, but for a short time.
 - Clients such as a mail server or a database server that back up large amounts of data in few transactions can cause small numbers of transactions that take a long time to complete. The transactions might use a small amount of space in the active log, but for a long time.
- o Network connection types
 - Backup operations that occur over fast network connections cause transactions that complete more quickly. The transactions use space in the active log for a shorter time.
 - Backup operations that occur over relatively slower connections cause transactions that take a longer time to complete. The transactions use space in the active log for a longer time.

If the server is handling transactions with a wide variety of characteristics, the space that is used for the active log might increase and decrease significantly over time. For such a server, you might need to ensure that the active log typically has a smaller percentage of its space used. The extra space allows the active log to grow for transactions that take a long time to complete.

2. Monitor the archive log to ensure that space is always available.

Remember: If the archive log becomes full, and the failover archive log becomes full, the active log can become full, and the server stops. The goal is to make enough space available to the archive log so that it never uses all its available space. You are likely to notice the following pattern:

- a. Initially, the archive log grows rapidly as typical client-backup operations occur.
- b. Database backups occur regularly, either as scheduled or done manually.
- c. After at least two full database backups occur, log pruning occurs automatically. The space that is used by the archive log decreases when the pruning occurs.
- d. Normal client operations continue, and the archive log grows again.
- e. Database backups occur regularly, and log pruning occurs as often as full database backups occur.

With this pattern, the archive log grows initially, decreases, and then might grow again. Over time, as normal operations continue, the amount of space that is used by the archive log should reach a relatively constant level.

If the archive log continues to grow, consider taking one or both of these actions:

- Add space to the archive log. You might need to move the archive log to a different file system.
 - Increase the frequency of full database backups, so that log pruning occurs more frequently.
3. If you defined a directory for the failover archive log, determine whether any logs get stored in that directory during normal operations. If the failover log space is being used, consider increasing the size of the archive log. The goal is that the failover archive log is used only under unusual conditions, not in normal operation.

Linux: Installing an IBM Spectrum Protect server fix pack

IBM Spectrum Protect™ maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.

Before you begin

To install a fix pack or interim fix to the server, install the server at the level on which you want to run it. You do not have to start the server installation at the base release level. For example, if you currently have V8.1.1 installed, you can go directly to the latest fix pack for V8.1. You do not have to start with the V8.1.0 installation if a maintenance update is available.

You must have the IBM Spectrum Protect license package installed. The license package is provided with the purchase of a base release. When you download a fix pack or interim fix from Fix Central, install the server license that is available on the Passport Advantage® website. To display messages and help in a language other than US English, install the language package of your choice.

If you upgrade the server to V8.1.5 or later, and then revert the server to a level that is earlier than V8.1.5, you must restore the database to a point in time before the upgrade. During the upgrade process, complete the required steps to ensure that the database can be restored: back up the database, the volume history file, the device configuration file, and the server options file. For more information, see [Linux: Reverting from Version 8.1.5 to a previous server](#).

If you are using the client management service, ensure that you upgrade it to the same version as the IBM Spectrum Protect server.

Ensure that you retain the installation media from the base release of the installed server. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Visit the IBM® Support Portal for the following information:

- A list of the latest maintenance and download fixes. Click **Downloads** and apply any applicable fixes.
- Details about obtaining a base license package. Search for **Downloads > Passport Advantage**.
- Supported platforms and system requirements. Search for **IBM Spectrum Protect supported operating systems**.

Ensure that you upgrade the server before you upgrade backup-archive clients. If you do not upgrade the server first, communication between the server and clients might be interrupted.

Attention: Do not alter the DB2® software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Procedure

To install a fix pack or interim fix, complete the following steps:

1. Back up the database. The preferred method is to use a snapshot backup. A snapshot backup is a full database backup that does not interrupt any scheduled database backups. For example, issue the following IBM Spectrum Protect administrative command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information. Issue the following IBM Spectrum Protect administrative command:

```
backup devconfig filenames=file_name
```

where *file_name* specifies the name of the file in which to store device configuration information.

3. Save the volume history file to another directory or rename the file. Issue the following IBM Spectrum Protect administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named `dsmserv.opt`. The file is in the server instance directory.
5. Halt the server before installing a fix pack or interim fix. Use the HALT command.
6. Ensure that extra space is available in the installation directory. The installation of this fix pack might require additional temporary disk space in the installation directory of the server. The amount of additional disk space can be as much as that required for installing a new database as part of an IBM Spectrum Protect installation. The IBM Spectrum Protect installation wizard displays the amount of space that is required for installing the fix pack and the available amount. If the required amount of space is greater than the available amount, the installation stops. If the installation stops, add the required disk space to the file system and restart the installation.
7. **Linux** Log in as the root user.
8. Obtain the package file for the fix pack or interim fix that you want to install from the IBM Support Portal, Passport Advantage, or Fix Central.
9. **Linux** Change to the directory where you placed the executable file and complete the following steps.
Tip: The files are extracted to the current directory. Ensure that the executable file is in the directory where you want the extracted files to be located.
 - a. Change file permissions by entering the following command:

```
chmod a+x 8.x.x.x-IBM-SPSRV-platform.bin
```

where *platform* denotes the architecture that IBM Spectrum Protect is to be installed on.

- b. Issue the following command to extract the installation files:

```
./8.x.x.x-IBM-SPSRV-platform.bin
```

10. Select one of the following ways of installing IBM Spectrum Protect.

Important: After a fix pack is installed, it is not necessary to go through the configuration again. You can stop after completing the installation, fix any errors, then restart your servers.

Install the IBM Spectrum Protect software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

Linux: Installing IBM Spectrum Protect by using the installation wizard

Tip: After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.

Command line in console mode

Follow the instructions for your operating system:

Linux: Installing IBM Spectrum Protect by using console mode

Silent mode

Follow the instructions for your operating system:

Linux: Installing IBM Spectrum Protect in silent mode

Tip: If you have multiple server instances on your system, run the installation wizard only once. The installation wizard upgrades all server instances.

Results

Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click File > View Log. To collect log files, from the IBM Installation Manager tool, click Help > Export Data for Problem Analysis.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

- **Linux** `/var/ibm/InstallationManager/logs`

Linux: Reverting from Version 8.1.5 to a previous server

If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the

preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect™ server with minimal loss of data.

Before you begin

You must have the following items from the earlier version of the server:

- Server database backup
- Volume history file
- Device configuration file
- Server options file

About this task

Use the same instructions whether you are reverting within releases or to an earlier release, for example, from 8.1.3 to 8.1.2 or from 8.1.3 to 7.1.2. The older version must match the version that you used before the upgrade to V8.1.

Attention: Specify the REUSEDELAY parameter to help prevent backup-archive client data loss when you revert the server to a previous version.

Steps for reverting to the previous server version

About this task

Complete the following steps on the system that has the V8.1 server.

Procedure

1. Halt the server to shut down all server operations by using the HALT command.
2. Remove the database from the database manager, then delete the database and recovery log directories.
 - a. Manually remove the database. One way to remove it is by issuing this command: **Linux**

```
dmserv removedb tsmdb1
```

- b. If you must reuse the space that is occupied by the database and recovery log directories, you can now delete these directories.
3. Use the uninstallation program to uninstall the V8.1 server. Uninstallation removes the server and the database manager, with their directories. For details, see Linux: Uninstalling IBM Spectrum Protect.
 4. Stop the cluster service. Reinstall the version of the server program that you were using before the upgrade to V8.1.5. This version must match the version that your server was running when you created the database backup that you restore in a later step. For example, the server was at V7.1.7 before the upgrade, and you intend to use the database backup that was in use on this server. You must install the V7.1.7 fix pack to be able to restore the database backup.
 5. Configure the new server database by using the configuration wizard. To start the wizard, issue the following command:

```
Linux  
. /dsmicfgx
```

6. Ensure that no servers are running in the background.
7. Restore the database to a point in time before the upgrade.
8. Copy the following files to the instance directory.
 - Device configuration file
 - Volume history file
 - The server options file (typically dmserv.opt)
9. If you enabled data deduplication for any FILE-type storage pools that existed before the upgrade, or if you moved data that existed before the upgrade into new storage pools while using the V8.1.5 server, you must complete additional recovery steps. For more details, see Additional recovery steps if you created new storage pools or enabled data deduplication.
10. If the REUSEDELAY parameter setting on storage pools is less than the age of the database that you restored, restore volumes on any sequential-access storage pools that were reclaimed after that database backup. Use the RESTORE VOLUME command.
If you do not have a backup of a storage pool, audit the reclaimed volumes by using the AUDIT VOLUME command, with the FIX=YES parameter to resolve inconsistencies. For example:

```
audit volume volume_name fix=yes
```

11. If client backup or archive operations were completed using the V8.1 server, audit the storage pool volumes on which the data was stored.

Additional recovery steps if you created new storage pools or enabled data deduplication

If you created new storage pools, turned on data deduplication for any FILE-type storage pools, or did both while your server was running as a V8.1.5 server, you must complete more steps to return to the previous server version.

Before you begin

To complete this task, you must have a complete backup of the storage pool that was created before the upgrade to V8.1.5.

About this task

Use this information if you did either or both of the following actions while your server was running as a V8.1.5 server:

- You enabled the data deduplication function for any storage pools that existed before the upgrade to V8.1.5 program. Data deduplication applies only to storage pools that use a FILE device type.
- You created new primary storage pools after the upgrade *and* moved data that was stored in other storage pools into the new storage pools.

Complete these steps after the server is again restored to V7.

Procedure

- For each storage pool for which you enabled the data deduplication function, restore the entire storage pool by using the RESTORE STGPPOOL command.
- For storage pools that you created after the upgrade, determine what action to take. Data that was moved from existing V8 storage pools into the new storage pools might be lost because the new storage pools no longer exist in your restored V8 server. Possible recovery depends on the type of storage pool:
 - If data was moved from V8 DISK-type storage pools into a new storage pool, space that was occupied by the data that was moved was probably reused. Therefore, you must restore the original V8 storage pools by using the storage pool backups that were created before the upgrade to V8.1.5.

If *no* data was moved from V8 DISK-type storage pools into a new storage pool, then audit the storage pool volumes in these DISK-type storage pools.
 - If data was moved from V8 sequential-access storage pools into a new storage pool, that data might still exist and be usable in storage pool volumes on the restored V8 server. The data might be usable if the REUSEDELAY parameter for the storage pool was set to a value that prevented reclamation while the server was running as a V8.1.5 server. If any volumes were reclaimed while the server was running as a V8.1.5 server, restore those volumes from storage pool backups that were created before the upgrade to V8.1.5.

Linux: Reference: DB2 commands for IBM Spectrum Protect server databases

Use this list as reference when you are directed to issue DB2® commands by IBM® support.

Purpose

After using the wizards to install and configure IBM Spectrum Protect™, you seldom need to issue DB2 commands. A limited set of DB2 commands that you might use or be asked to issue are listed in Table 1. This list is supplemental material only and is not a comprehensive list. There is no implication that an IBM Spectrum Protect administrator will use it on a daily or ongoing basis. Samples of some commands are provided. Details of output are not listed.

For a full explanation of the commands described here and of their syntax, see the DB2 product information.

Table 1. DB2 commands

Command	Description	Example
---------	-------------	---------

Command	Description	Example
db2icrt	<p>Creates DB2 instances in the home directory of the instance owner.</p> <p>Tip: The IBM Spectrum Protect configuration wizard creates the instance used by the server and database. After a server is installed and configured through the configuration wizard, the db2icrt command is generally not used.</p> <p>Linux This utility is in the DB2DIR/instance directory, where DB2DIR represents the installation location where the current version of the DB2 database system is installed.</p>	<p>Manually create an IBM Spectrum Protect instance. Enter the command on one line:</p> <pre>/opt/tivoli /tsm/db2/in stance/ db2icrt -a server -u instance_na me instance_na me</pre>
db2set	Displays DB2 variables.	<p>List DB2 variables:</p> <pre>db2set</pre>
CATALOG DATABASE	Stores database location information in the system database directory. The database can be located either on the local workstation or on a remote database partition server. The server configuration wizard takes care of any catalog needed for using the server database. Run this command manually, after a server is configured and running, only if something in the environment changes or is damaged.	<p>Catalog the database:</p> <pre>db2 catalog database tsmdb1</pre>
CONNECT TO DATABASE	Connects to a specified database for command-line interface (CLI) use.	<p>Connect to the IBM Spectrum Protect database from a DB2 CLI:</p> <pre>db2 connect to tsmdb1</pre>
GET DATABASE CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	<p>Show the configuration information for a database alias:</p> <pre>db2 get db cfg for tsmdb1</pre> <p>Retrieve information in order to verify settings such as database configuration, log mode, and maintenance.</p> <pre>db2 get db config for tsmdb1 show detail</pre>

Command	Description	Example
GET DATA BASE MAN AGE R CON FIGU RATI ON	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Retrieve configuration information for the database manager: db2 get dbm cfg
GET HEAL TH SNA PSH OT	Retrieves the health status information for the database manager and its databases. The information returned represents a snapshot of the health state at the time the command was issued. IBM Spectrum Protect monitors the state of the database using the health snapshot and other mechanisms that are provided by DB2. There might be cases where the health snapshot or other DB2 documentation indicates that an item or database resource might be in an alert state. Such a case indicates that action must be considered to remedy the situation. IBM Spectrum Protect monitors the condition and responds appropriately. Not all declared alerts by the DB2 database are acted on.	Receive a report on DB2 health monitor indicators: db2 get health snapshot for database on tsmdb1
GRA NT (Data base Auth oritie s)	Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.	Grant access to the user ID itmuser: db2 GRANT CONNECT ON DATABASE TO USER itmuser db2 GRANT CREATETAB ON DATABASE TO USER itmuser
RUN STAT S	Updates statistics about the characteristics of a table and associated indexes or statistical views. These characteristics include number of records, number of pages, and average record length. To see a table, issue this utility after updating or reorganizing the table. A view must be enabled for optimization before its statistics can be used to optimize a query. A view that is enabled for optimization is known as a statistical view. Use the DB2 ALTER VIEW statement to enable a view for optimization. Issue the RUNSTATS utility when changes to underlying tables substantially affect the rows returned by the view. Tip: The server configures DB2 to run the RUNSTATS command as needed.	Update statistics on a single table. db2 runstats on table SCHEMA_NAME .TABLE_NAME with distributio n and sampled detailed indexes all
SET SCH EMA	Changes the value of the CURRENT SCHEMA special register, in preparation for issuing SQL commands directly through the DB2 CLI. Tip: A special register is a storage area that is defined for an application process by the database manager. It is used to store information that can be referenced in SQL statements.	Set the schema for IBM Spectrum Protect: db2 set schema tsmdb1

Command	Description	Example
START DATABASE MANAGER	Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts. Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Start the database manager: <code>db2start</code>
STOP DATABASE MANAGER	Stops the current database manager instance. Unless explicitly stopped, the database manager continues to be active. This command does not stop the database manager instance if any applications are connected to databases. If there are no database connections, but there are instance attachments, the command forces the instance attachments to stop first. Then, it stops the database manager. This command also deactivates any outstanding database activations before stopping the database manager. This command is not valid on a client. The server starts and stops the instance and database whenever the server starts and halts. Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Stop the database manager: <code>db2 stop dbm</code>

Linux: Uninstalling IBM Spectrum Protect

You can use the following procedures to uninstall IBM Spectrum Protect™. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Before you begin

Complete the following steps before you uninstall IBM Spectrum Protect:

- Complete a full database backup.
- Save a copy of the volume history and device configuration files.
- Store the output volumes in a safe location.

About this task

You can uninstall IBM Spectrum Protect by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

- **Linux: Uninstalling IBM Spectrum Protect by using a graphical wizard**
You can uninstall IBM Spectrum Protect by using the IBM® Installation Manager installation wizard.
- **Linux: Uninstalling IBM Spectrum Protect in console mode**
To uninstall IBM Spectrum Protect by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.
- **Linux: Uninstalling IBM Spectrum Protect in silent mode**
To uninstall IBM Spectrum Protect in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.
- **Linux: Uninstalling and reinstalling IBM Spectrum Protect**
If you plan to manually reinstall IBM Spectrum Protect instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.
- **Linux: Uninstalling IBM Installation Manager**
You can uninstall IBM Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

What to do next

See Linux: Installing the server components for installation steps to reinstall the IBM Spectrum Protect components.

Linux: Uninstalling IBM Spectrum Protect by using a graphical wizard

You can uninstall IBM Spectrum Protect™ by using the IBM® Installation Manager installation wizard.

Procedure

1. Start the Installation Manager.

Linux In the directory where the Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command:

```
./IBMIM
```

2. Click Uninstall.
3. Select IBM Spectrum Protect server, and click Next.
4. Click Uninstall.
5. Click Finish.

Linux: Uninstalling IBM Spectrum Protect in console mode

To uninstall IBM Spectrum Protect™ by using the command line, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **Linux** eclipse/tools

For example:

- o **Linux** /opt/IBM/InstallationManager/eclipse/tools

2. From the tools directory, issue the following command:

- o **Linux** ./imcl -c

3. To uninstall, enter 5.
4. Choose to uninstall from the IBM Spectrum Protect package group.
5. Enter N for Next.
6. Choose to uninstall the IBM Spectrum Protect server package.
7. Enter N for Next.
8. Enter U for Uninstall.
9. Enter F for Finish.

Linux: Uninstalling IBM Spectrum Protect in silent mode

To uninstall IBM Spectrum Protect™ in silent mode, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the IBM Spectrum Protect server components. IBM Spectrum Protect includes a sample response file, `uninstall_response_sample.xml`, in the input directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

If you want to uninstall all IBM Spectrum Protect components, leave `modify="false"` set for each component in the response file. If you do not want to uninstall a component, set the value to `modify="true"`.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **Linux** eclipse/tools

For example:

- o `Linux` /opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command, where *response_file* represents the response file path, including the file name:

```
Linux  
./imcl -input response_file -silent
```

The following command is an example:

```
Linux  
./imcl -input /tmp/input/uninstall_response.xml -silent
```

Linux: Uninstalling and reinstalling IBM Spectrum Protect

If you plan to manually reinstall IBM Spectrum Protect™ instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.

About this task

To manually uninstall and reinstall IBM Spectrum Protect, complete the following steps:

1. `Linux` Make a list of your current server instances before proceeding to the uninstallation. Run the following command:

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

2. Run the following commands for every server instance:

```
Linux  
db2 attach to instance_name  
db2 get dbm cfg show detail  
db2 detach
```

Keep a record of the database path for each instance.

3. Uninstall IBM Spectrum Protect. See Linux: Uninstalling IBM Spectrum Protect.
4. When you uninstall any supported version of IBM Spectrum Protect, including a fix pack, an instance file is created. The instance file is created to help reinstall IBM Spectrum Protect. Check this file and use the information when you are prompted for the instance credentials when reinstalling. In silent installation mode, you provide these credentials using the `INSTANCE_CRED` variable.

You can find the instance file in the following location:

- o `Linux` /etc/tivoli/tsm/instanceList.obj
5. Reinstall IBM Spectrum Protect. See Linux: Installing the server components.

If the `instanceList.obj` file does not exist, you need to recreate your server instances using the following steps:

- a. Recreate your server instances. See Linux: Creating the server instance.
Tip: The installation wizard configures the server instances but you must verify that they exist. If they do not exist, you must manually configure them.
- b. Catalog the database. Log in to each server instance as the instance user, one at a time, and issue the following commands:

```
Linux  
db2 catalog database tsmdb1  
db2 attach to instance_name  
db2 update dbm cfg using dftdbpath instance_directory  
db2 detach
```

- c. `Linux` Verify that the server instance was created successfully. Issue this command:

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

- d. Verify that IBM Spectrum Protect recognizes the server instance by listing your directories. Your home directory appears if you did not change it. Your instance directory does appear if you used the configuration wizard. Issue this command:

```
db2 list database directory
```

If you see TSMDDB1 listed, you can start the server.

Linux: Uninstalling IBM Installation Manager

You can uninstall IBM® Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

Before you begin

Before you uninstall IBM Installation Manager, you must ensure that all packages that were installed by IBM Installation Manager are uninstalled. Close IBM Installation Manager before you start the uninstall process.

Linux To view installed packages, issue the following command from a command line:

```
cd /opt/IBM/InstallationManager/eclipse/tools
./imcl listInstalledPackages
```

Procedure

To uninstall IBM Installation Manager, complete the following steps:

Linux

1. Open a command line and change directories to `/var/ibm/InstallationManager/uninstall`.
2. Issue the following command:

```
./uninstall
```

Restriction: You must be logged in to the system as the `root` user ID.

Windows: Installing the server

Installation of the server includes planning, installation, and initial configuration.

- **Windows: Planning to install the server**
Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.
- **Windows: Installing the server components**
To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.
- **Windows: Taking the first steps after you install IBM Spectrum Protect**
After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect instance.
- **Windows: Installing an IBM Spectrum Protect server fix pack**
IBM Spectrum Protect maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.
- **Windows: Reverting from Version 8.1.5 to a previous server**
If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect server with minimal loss of data.
- **Windows: Reference: DB2 commands for IBM Spectrum Protect server databases**
Use this list as reference when you are directed to issue DB2® commands by IBM® support.
- **Windows: Uninstalling IBM Spectrum Protect**
You can use the following procedures to uninstall IBM Spectrum Protect. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Windows: Planning to install the server

Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect™ server-managed storage.

- **Windows: What you should know first**
Before installing IBM Spectrum Protect, be familiar with your operating systems, storage devices, communication protocols, and system configurations.
- **Windows: Planning for optimal performance**
Before you install the IBM Spectrum Protect server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.
- **Windows** **Windows: Minimum system requirements for Windows systems**
Before you install an IBM Spectrum Protect server on a Windows operating system, review the hardware and software requirements.
- **Windows: IBM Installation Manager**
IBM Spectrum Protect uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.
- **Windows: Worksheets for planning details for the server**
You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect server. You can also use them to keep track of names and user IDs.
- **Windows: Capacity planning**
Capacity planning for IBM Spectrum Protect includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.
- **Windows: Server naming best practices**
Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect server.
- **Windows: Installation directories**
Installation directories for the IBM Spectrum Protect server include the server, DB2®, device, language, and other directories. Each one contains several additional directories.

Windows: What you should know first

Before installing IBM Spectrum Protect™, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

Server maintenance releases, client software, and publications are available from the IBM® Support Portal.

Windows **Restriction:** You cannot install and run the Version 8.1.5 server on a system that already has DB2® installed on it, whether DB2 was installed by itself or as part of some other application. The V8.1.5 server requires the installation and use of the DB2 version that is packaged with the V8.1.5 server. No other version of DB2 can exist on the system.

Windows You can install the IBM Spectrum Protect server on a domain controller. The server can have heavy processor usage, however, and that might affect and stall other applications.

Experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. Do not, however, use DB2 tools to change DB2 configuration settings from those that are preset by IBM Spectrum Protect, or alter the DB2 environment for IBM Spectrum Protect in other ways, such as with other products. The V8.1.5 server has been built and tested extensively using the data definition language (DDL) and database configuration that the server deploys.

Attention: Do not alter the DB2 software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Windows: Planning for optimal performance

Before you install the IBM Spectrum Protect™ server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.

Procedure

1. Review **Windows: What you should know first**.
2. Review each of the following sub-sections.
 - **Windows: Planning for the server hardware and the operating system**
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

- **Windows: Planning for the server database disks**
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
- **Windows: Planning for the server recovery log disks**
Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.
- **Windows: Planning for directory-container and cloud-container storage pools**
Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.
- **Windows: Planning for storage pools in DISK or FILE device classes**
Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.
- **Windows: Planning for the correct type of storage technology**
Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect.
- **Windows: Applying best practices to the server installation**
Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Windows: Planning for the server hardware and the operating system

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Does the operating system and hardware meet or exceed requirements?</p> <ul style="list-style-type: none"> • Number and speed of processors • System memory • Supported operating system level 	<p>If you are using the minimum required amount of memory, you can support a minimal workload.</p> <p>You can experiment by adding more system memory to determine whether the performance is improved. Then, decide whether you want to keep the system memory dedicated to the server. Test the memory variations by using the entire daily cycle of the server workload.</p> <p>If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.</p>	<p>Review operating system requirements at technote 1243309.</p> <p>Additionally, review the guidance in Tuning tasks for operating systems and other applications.</p> <p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For more information about sizing requirements for the server and storage, see the IBM Spectrum Protect™ Blueprint.</p>
<p>Are disks configured for optimal performance?</p>	<p>The amount of tuning that can be done for different disk systems varies. Ensure that the appropriate queue depths and other disk system options are set.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"

Question	Tasks, characteristics, options, or settings	More information
Does the server have enough memory?	<p>Heavier workloads and advanced features such as data deduplication and node replication require more than the minimum system memory that is specified in the system requirements document.</p> <p>For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements:</p> <ul style="list-style-type: none"> • For databases less than 500 GB, you need 16 GB of memory. • For databases with a size of 500 GB - 1 TB, you need 24 GB of memory. • For databases with a size of 1 TB - 1.5 TB, you need 32 GB of memory. • For databases greater than 1.5 TB, you need 40 GB of memory. <p>Ensure that you allocate extra space for the active log and the archive log for replication processing.</p>	<p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication • Memory requirements
Does the system have enough host bus adapters (HBAs) to handle the data operations that the IBM Spectrum Protect server must run simultaneously?	<p>Understand what operations require use of HBAs at the same time.</p> <p>For example, a server must store 1 GB/sec of backup data while also doing storage pool migration that requires 0.5 GB/sec capacity to complete. The HBAs must be able to handle all of the data at the speed required.</p>	See Tuning HBA capacity.
Is network bandwidth greater than the planned maximum throughput for backups?	<p>Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments.</p> <p>For node replication, network bandwidth must be greater than the planned maximum throughput.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Tuning network performance • Checklist for node replication

Question	Tasks, characteristics, options, or settings	More information
Are you using a preferred file system for IBM Spectrum Protect server files?	Use a file system that ensures optimal performance and data availability. The server uses direct I/O with file systems that support the feature. Using direct I/O can improve throughput and reduce processor use. For more information about the preferred file system for your operating system, see IBM Spectrum Protect server-supported file systems.	For more information, see Configuring the operating system for disk performance.
Are you planning to configure enough paging space?	<p>Paging space, or swap space, extends the memory that is available for processing. When the amount of free RAM in the system is low, programs or data that is not in use are moved from memory to paging space. This action releases memory for other activities, such as database operations.</p> <p>Windows Paging space is automatically configured.</p>	

Windows: Planning for the server database disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
Is the database on fast, low-latency disks?	<p>Do not use the following drives for the IBM Spectrum Protect™ database:</p> <ul style="list-style-type: none"> • Nearline SAS (NL-SAS) • Serial Advanced Technology Attachment (SATA) • Parallel Advanced Technology Attachment (PATA) <p>Do not use internal disks that are included by default in most server hardware.</p> <p>Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.</p> <p>If you plan to use the data deduplication functions of IBM Spectrum Protect, focus on disk performance in terms of I/O operations per second (IOPS).</p>	For more information, see Checklist for data deduplication.
Is the database stored on disks or LUNs that are separate from disks or LUNs that are used for the active log, archive log, and storage pool volumes?	<p>Separation of the server database from other server components helps reduce contention for the same resources by different operations that must run at the same time.</p> <p>Tip: The database and the archive log can share an array when you use solid-state drive (SSD) technology.</p>	
If you are using RAID, do you know how to select the optimal RAID level for your system? Are you defining all LUNs with the same size and type of RAID?	<p>When a system must do large numbers of writes, RAID 10 outperforms RAID 5. However, RAID 10 requires more disks than RAID 5 for the same amount of usable storage.</p> <p>If your disk system is RAID, define all your LUNs with the same size and type of RAID. For example, do not mix 4+1 RAID 5 with 4+2 RAID 6.</p>	
If an option to set the strip size or segment size is available, are you planning to optimize the size when you configure the disk system?	If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.	The block size that is used for the database varies depending on the table space. Most table spaces use 8 KB blocks, but some use 32 KB blocks.

Question	Tasks, characteristics, options, or settings	More information
<p>Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database?</p> <p>Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.</p>	<p>Heavier workloads and use of some features require more database storage paths than the minimum requirements.</p> <p>Server operations such as data deduplication drive a high number of input/output operations per second (IOPS) for the database. Such operations perform better when the database has more directories.</p> <p>For server databases that are larger than 2 TB or are expected to grow to that size, use eight directories.</p> <p>Consider planned growth of the system when you determine how many storage paths to create. The server uses the higher number of storage paths more effectively if the storage paths are present when the server is first created.</p> <p>Use the <i>DB2_PARALLEL_IO</i> variable to force parallel I/O to occur on table spaces that have one container, or on table spaces that have containers on more than one physical disk. If you do not set the <i>DB2_PARALLEL_IO</i> variable, I/O parallelism is equal to the number of containers that are used by the table space. For example, if a table space spans four containers, the level of I/O parallelism that is used is 4.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For help with forecasting growth when the server deduplicates data, see technote 1596944.</p> <p>For the most recent information about database size, database reorganization, and performance considerations for IBM Spectrum Protect servers, see technote 1683633.</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>
<p>Are all directories for the database the same size?</p>	<p>Directories that are all the same size ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching.</p> <p>This guideline also applies if you must add storage paths after the initial configuration of the server.</p>	
<p>Are you planning to raise the queue depth of the database LUNs on AIX® systems?</p>	<p>The default queue depth is often too low.</p>	<p>See Configuring AIX systems for disk performance.</p>

Windows: Planning for the server recovery log disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Are the active log and archive log stored on disks or LUNs that are separate from what is used for the database and storage pool volumes?</p>	<p>Ensure that the disks where you place the active log are not used for other server or system purposes. Do not place the active log on disks that contain the server database, the archive log, or system files such as page or swap space.</p>	<p>Separation of the server database, active log, and archive log helps to reduce contention for the same resources by different operations that must run at the same time.</p>

Question	Tasks, characteristics, options, or settings	More information
Are the logs on disks that have nonvolatile write cache?	Nonvolatile write cache allows data to be written to the logs as fast as possible. Faster write operations for the logs can improve performance for server operations.	
Are you setting the logs to a size that adequately supports the workload?	<p>If you are not sure about the workload, use the largest size that you can.</p> <p>Active log The maximum size is 512 GB, set with the ACTIVELOGSIZE server option.</p> <p>Ensure that there is at least 8 GB of free space on the active log file system after the fixed size active logs are created.</p> <p>Archive log The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Make the archive log at least as large as the active log.</p>	<ul style="list-style-type: none"> For log sizing details, see the recovery log information in technote 1421060. For information about sizing when you use data deduplication, see Checklist for data deduplication.
Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?	The archive failover log is for emergency use by the server when the archive log becomes full. Slower disks can be used for the archive failover log.	<p>Use the ARCHFAILOVERLOGDIRECTORY server option to specify the location of the archive failover log.</p> <p>Monitor the usage of the directory for the archive failover log. If the archive failover log must be used by the server, the space for the archive log might not be large enough.</p>
If you are mirroring the active log, are you using only one type of mirroring?	<p>You can mirror the log by using one of the following methods. Use only one type of mirroring for the log.</p> <ul style="list-style-type: none"> Use the MIRRORLOGDIRECTORY option that is available for the IBM Spectrum Protect™ server to specify a mirror location. Use software mirroring, such as Logical Volume Manager (LVM) on AIX®. Use mirroring in the disk system hardware. 	<p>If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability.</p> <p>For more information, see Configuring and tuning the recovery log.</p>

Windows: Planning for directory-container and cloud-container storage pools

Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.

Question	Tasks, characteristics, options, or settings	More information
----------	--	------------------

Question	Tasks, characteristics, options, or settings	More information
<p>Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the IBM Spectrum Protect™ database?</p>	<p>Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing.</p> <p>Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication), add 1000 IOPS to this minimum.</p> <p>For example, an IBM Spectrum Protect server that is ingesting 3 TB of data per day would need 6000 IOPS for the database disks:</p> $3000 \text{ IOPS minimum} + 3000 (3 \text{ TB} \times 1000 \text{ IOPS}) = 6000 \text{ IOPS}$	<p>For recommendations about disk selection, see "Planning for server database disks".</p> <p>For more information about IOPS, see the IBM Spectrum Protect Blueprints.</p>
<p>Do you have enough memory for the size of your database?</p>	<p>Use a minimum of 40 GB of system memory for IBM Spectrum Protect servers, with a database size of 100 GB, that are deduplicating data. If the retained capacity of backup data grows, the memory requirement might need to be higher.</p> <p>Monitor memory usage regularly to determine whether more memory is required.</p> <p>Use more system memory to improve caching of database pages. The following memory size guidelines are based on the daily amount of new data that you back up:</p> <ul style="list-style-type: none"> • 128 GB of system memory for daily backups of data, where the database size is 1 - 2 TB • 192 GB of system memory for daily backups of data, where the database size is 2 - 4 TB 	<p>Memory requirements</p>

Question	Tasks, characteristics, options, or settings	More information
<p>Have you properly sized the storage capacity for the database active log and archive log?</p>	<p>Configure the server to have a minimum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072.</p> <p>The suggested starting size for the archive log is 1 TB. The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Ensure that there is at least 10% extra disk space for the file system than the size of the archive log.</p> <p>Use a directory for the database archive logs with an initial free capacity of at least 1 TB. Specify the directory by using the ARCHLOGDIRECTORY server option.</p> <p>Define space for the archive failover log by using the ARCHFAILOVERLOGDIRECTORY server option.</p>	<p>For more information about sizing for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Is compression enabled for the archive log and database backups?</p>	<p>Enable the ARCHLOGCOMPRESS server option to save storage space.</p> <p>This compression option is different from inline compression. Inline compression is enabled by default with IBM Spectrum Protect V7.1.5 and later.</p> <p>Restriction: Do not use this option if the amount of backed up data exceeds 6 TB per day.</p>	<p>For more information about compression for your system, see the IBM Spectrum Protect Blueprints.</p>
<p>Are the IBM Spectrum Protect database and logs on separate disk volumes (LUNs)?</p> <p>Is the disk that is used for the database configured according to best practices for a transactional database?</p>	<p>The database must not share disk volumes with IBM Spectrum Protect database logs or storage pools, or with any other application or file system.</p>	<p>For more information about server database and recovery log configuration, see Server database and recovery log configuration and tuning.</p>
<p>Are you using a minimum of eight (2.2 GHz or equivalent) processor cores for each IBM Spectrum Protect server that you plan to use with data deduplication?</p>	<p>If you are planning to use client-side data deduplication, verify that client systems have adequate resources available during a backup operation to complete data deduplication processing. Use a processor that is at least the minimum equivalent of one 2.2 GHz processor core per backup process with client-side data deduplication.</p>	<ul style="list-style-type: none"> • Effective planning and use of deduplication • IBM Spectrum Protect Blueprints

Question	Tasks, characteristics, options, or settings	More information
Did you allocate enough storage space for the database?	<p>For a rough estimate, plan for 100 GB of database storage for every 50 TB of data that is to be protected in deduplicated storage pools. <i>Protected data</i> is the amount of data before data deduplication, including all versions of objects stored.</p> <p>As a best practice, define a new container storage pool exclusively for data deduplication. Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.</p>	
Have you estimated storage pool capacity to configure enough space for the size of your environment?	<p>You can estimate capacity requirements for a deduplicated storage pool by using the following technique:</p> <ol style="list-style-type: none"> 1. Estimate the base size of the source data. 2. Estimate the daily backup size by using an estimated change and growth rate. 3. Determine retention requirements. 4. Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements. 5. Apply the deduplication ratio factor. 6. Apply the compression ratio factor. 7. Round up the estimate to consider transient storage pool usage. 	For an example of using this technique, see Effective planning and use of deduplication.
Have you distributed disk I/O over many disk devices and controllers?	<p>Use arrays that consist of as many disks as possible, which is sometimes referred to as wide striping. Ensure that you use one database directory per distinct array on the subsystem.</p> <p>Set the <i>DB2_PARALLEL_IO</i> registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks.</p> <p>When I/O bandwidth is available and the files are large, for example 1 MB, the process of finding duplicates can occupy the resources of an entire processor. When files are smaller, other bottlenecks can occur.</p> <p>Specify eight or more file systems for the deduplicated storage pool device class so that I/O is distributed across as many LUNs and physical devices as possible.</p>	<p>For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes".</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>

Question	Tasks, characteristics, options, or settings	More information
Have you scheduled daily operations based on your backup strategy?	<p>The best practice sequence of operations is in the following order:</p> <ol style="list-style-type: none"> 1. Client backup 2. Storage pool protection 3. Node replication 4. Database backup 5. Expire inventory 	<ul style="list-style-type: none"> • Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools
Do you have enough storage to manage the DB2® lock list?	<p>If you deduplicate data that includes large files or large numbers of files concurrently, the process can result in insufficient storage space. When the lock list storage is insufficient, backup failures, data management process failures, or server outages can occur.</p> <p>File sizes greater than 500 GB that are processed by data deduplication are most likely to deplete storage space. However, if many backup operations use client-side data deduplication, this problem can also occur with smaller-sized files.</p>	For information about tuning the DB2 LOCKLIST parameter, see Tuning server-side data deduplication.
Is sufficient bandwidth available to transfer data to an IBM Spectrum Protect server?	<p>To transfer data to an IBM Spectrum Protect server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required.</p> <p>Use a V7.1.5 server or higher to use inline compression and use a V7.1.6 or later client to enable enhanced compression processing.</p>	For more information, see the enablededup client option.
Have you determined how many storage pool directories to assign to each storage pool?	<p>Assign directories to a storage pool by using the DEFINE STGPOOLDIRECTORY command.</p> <p>Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).</p>	

Question	Tasks, characteristics, options, or settings	More information
<p>Did you allocate enough disk space in the cloud-container storage pool?</p>	<p>To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space:</p> <ul style="list-style-type: none"> • For serial-attached SCSI (SAS) and spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount, in terabytes, for disk space. • Provide 3 TB for flash-based storage systems with fast network connections to on-premises, high-performance cloud systems. • Provide 5 TB for solid-state drive (SSD) systems with fast network connections to high-performance cloud systems. 	
<p>Did you select the appropriate type of local storage?</p>	<p>Ensure that data transfers from local storage to cloud finish before the next backup cycle starts. Tip: Data is removed from local storage soon after it moves to the cloud. Use the following guidelines:</p> <ul style="list-style-type: none"> • Use flash or SSD for large systems that have high-performing cloud systems. Ensure that you have a dedicated 10 GB wide area network (WAN) link with a high-speed connection to the object storage. For example, use flash or SSD if you have a dedicated 10 GB WAN link plus a high-speed connection to either an IBM® Cloud Object Storage location or to an Amazon Simple Storage Service (Amazon S3) data center. • Use larger capacity 15000 rpm SAS disks for these scenarios: <ul style="list-style-type: none"> ◦ Medium-sized systems ◦ Slower cloud connections, for example, 1 GB ◦ When you use IBM Cloud Object Storage as your service provider across several regions • For SAS or spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount for disk space, in terabytes. 	

Windows: Planning for storage pools in DISK or FILE device classes

Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.

Question	Tasks, characteristics, options, or settings	More information
Can the storage pool LUNs sustain throughput rates for 256 KB sequential reads and writes to adequately handle the workload within the time constraints?	<p>When you are planning for peak loads, consider all the data that you want the server to read or write to the disk storage pools simultaneously. For example, consider the peak flow of data from client backup operations and server data-movement operations such as migration that run at the same time.</p> <p>The IBM Spectrum Protect™ server reads and writes to storage pools predominantly in 256 KB blocks.</p> <p>If the disk system includes the capability, configure the disk system for optimal performance with sequential read/write operations rather than random read/write operations.</p>	For more information, see Analyzing the basic performance of disk systems.
Is the disk configured to use read and write cache?	Use more cache for better performance.	
For storage pools that use FILE device classes, have you determined a good size to use for the storage pool volumes?	Review the information in Optimal number and size of volumes for storage pools that use disk. If you do not have the information to estimate a size for FILE device class volumes, start with volumes that are 50 GB.	Typically, problems arise more frequently when the volumes are too small. Few problems are reported when volumes are larger than needed. When you determine the volume size to use, as a precaution choose a size that might be larger than necessary.
For storage pools that use FILE device classes, are you using preallocated volumes?	<p>Scratch volumes can cause file fragmentation.</p> <p>To ensure that a storage pool does not run out of volumes, set the MAXSCRATCH parameter to a value greater than zero.</p>	<p>Use the DEFINE VOLUME server command to preallocate volumes in the storage pool.</p> <p>Use the DEFINE STGPOOL or UPDATE STGPOOL server command to set the MAXSCRATCH parameter.</p>
For storage pools that use FILE device classes, have you compared the maximum number of client sessions to the number of volumes that are defined?	Always maintain enough usable volumes in the storage pools to allow for the expected peak number of client sessions that run at one time. The volumes might be scratch volumes, empty volumes, or partly filled volumes.	For storage pools that use FILE device classes, only one session or process can write to a volume at the same time.

Question	Tasks, characteristics, options, or settings	More information
<p>For storage pools that use FILE device classes, have you set the MOUNTLIMIT parameter of the device class to a value that is high enough to account for the number of volumes that might be mounted in parallel?</p>	<p>For storage pools that use data deduplication, the MOUNTLIMIT parameter is typically in the range of 500 - 1000.</p> <p>Set the value for MOUNTLIMIT to the maximum number of mount points that are needed for all active sessions. Consider parameters that affect the maximum number of mount points that are needed:</p> <ul style="list-style-type: none"> • The MAXSESSIONS server option, which is the maximum number of IBM Spectrum Protect sessions that can run concurrently. • The MAXNUMMP parameter, which sets the maximum number of mount points that each client node can use. <p>For example, if the maximum number of client node backup sessions is typically 100 and each of the nodes has MAXNUMMP=2, multiply 100 nodes by the 2 mount points for each node to get the value of 200 for the MOUNTLIMIT parameter.</p>	<p>Use the REGISTER NODE or UPDATE NODE server command to set the MAXNUMMP parameter for client nodes.</p>
<p>For storage pools that use DISK device classes, have you determined how many storage pool volumes to put on each file system?</p>	<p>How you configure the storage for a storage pool that uses a DISK device class depends on whether you are using RAID for the disk system.</p> <p>If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system.</p> <p>If you are using RAID 5 with $n + 1$ volumes, configure the storage in one of the following ways:</p> <ul style="list-style-type: none"> • Configure n file systems on the LUN and define one storage pool volume per file system. • Configure one file system and n storage pool volumes for the LUN. 	<p>For an example layout that follows this guideline, see Sample layout of server storage pools.</p>
<p>Did you create your storage pools to distribute I/O across multiple file systems?</p>	<p>Ensure that each file system is on a different LUN on the disk system.</p> <p>Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.</p>	<p>For details, see the following topics:</p> <ul style="list-style-type: none"> • Tuning disk storage for the server • Tuning and configuring storage pools and volumes

Windows: Planning for the correct type of storage technology

Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect™.

Procedure

Review the following table to help you to choose the correct type of storage technology for the storage resources that the server requires.

Table 1. Storage technology types for IBM Spectrum Protect storage requirements

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
Solid-state disk (SSD)	Place the database on SSD in the following circumstances: <ul style="list-style-type: none"> You are using IBM Spectrum Protect data deduplication. You are backing up more than 8 TB of new data daily. 	If you place the IBM Spectrum Protect database on an SSD, as a best practice, place the active log on an SSD. If space is not available, use high-performance disk instead.	Save SSDs for use with the database and active log. The archive log and archive failover logs can be placed on slower storage technology types.	Save SSDs for use with the database and active log. Storage pools can be placed on slower storage technology types.
High-performance disk with the following characteristics: <ul style="list-style-type: none"> 15k rpm disk Fibre Channel or serial-attached SCSI (SAS) interface 	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. Isolate the server database from its logs and storage pools, and from data for other applications.	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use high-performance disks for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use high-performance disks for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.
Medium-performance or high-performance disk with the following characteristics: <ul style="list-style-type: none"> 10k rpm disk Fibre Channel or SAS interface 	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. Isolate the server database from its logs and storage pools, and from data for other applications.	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use medium-performance or high-performance disk for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use medium-performance or high-performance disk for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
SATA, network-attached storage	Do not use this storage for the database. Do not place the database on XIV storage systems.	Do not use this storage for the active log.	Use of this slower storage technology is acceptable because these logs are written once and infrequently read.	Use this slower storage technology in the following circumstances: <ul style="list-style-type: none"> • Data is infrequently written, for example written once. • Data is infrequently read.
Tape and virtual tape				Use for long-term retention or if data is infrequently used.

Windows: Applying best practices to the server installation

Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect™ solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Procedure

- The following best practices are the most important for optimal performance and problem prevention.
- Review the table to determine the best practices that apply to your environment.

Best practice	More information
Use fast disks for the server database. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.	Use fast, low-latency disks for the database. Using SSD is essential if you are using data deduplication and node replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: <ul style="list-style-type: none"> ◦ "Planning for server database disks" ◦ "Planning for the correct type of storage technology"
Ensure that the server system has enough memory.	Review operating system requirements in technote 1243309. Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and node replication can require more than the minimum memory that is specified in the system requirements document. If you plan to run multiple instances, each instance requires the memory that is listed for one server. Multiply the memory for one server by the number of instances that are planned for the system.

Best practice	More information
Separate the server database, the active log, the archive log, and disk storage pools from each other.	<p>Keep all IBM Spectrum Protect storage resources on separate disks. Keep storage pool disks separate from the disks for the server database and logs. Storage pool operations can interfere with database operations when both are on the same disks. Ideally, the server database and logs are also separated from each other. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> o "Planning for server database disks" o "Planning for server recovery log disks" o "Planning for storage pools in DISK or FILE device classes"
Use at least four directories for the server database. For larger servers or servers that use advanced features, use eight directories.	<p>Place each directory on a LUN that is isolated from other LUNs and from other applications.</p> <p>A server is considered to be large if its database is larger than 2 TB or is expected to grow to that size. Use eight directories for such servers.</p> <p>See "Planning for server database disks".</p>
If you are using data deduplication, node replication, or both, follow the guidelines for database configuration and other items.	<p>Configure the server database according to the guidelines, because the database is extremely important to how well the server runs when these features are being used. For details and more tips, see the following topics:</p> <ul style="list-style-type: none"> o Checklist for data deduplication o Checklist for node replication
For storage pools that use FILE type device classes, follow the guidelines for the size of storage pool volumes. Typically, 50 GB volumes are best.	<p>Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size.</p> <p>Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements.</p> <p>Isolate the storage devices that are used by IBM Spectrum Protect from other applications that have high I/O, and ensure that there is enough throughput to that storage.</p> <p>For more details, see Checklist for storage pools on DISK or FILE.</p>
Schedule IBM Spectrum Protect client operations and server maintenance activities to avoid or minimize overlap of operations.	<p>For more details, see the following topics:</p> <ul style="list-style-type: none"> o Tuning the schedule for daily operations o Checklist for server configuration
Monitor operations constantly.	<p>By monitoring, you can find problems early and more easily identify causes. Keep records of monitoring reports for up to a year to help you identify trends and plan for growth. See Monitoring and maintaining the environment for performance.</p>

Windows: Minimum system requirements for Windows systems

Before you install an IBM Spectrum Protect™ server on a Windows operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Spectrum Protect server installation

The optimal IBM Spectrum Protect environment is set up with data deduplication by using the IBM Spectrum Protect Blueprints.

For the most current information about IBM Spectrum Protect system requirements, see technote 1243309.

Windows: IBM Installation Manager

IBM Spectrum Protect™ uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install IBM Spectrum Protect. It must remain installed on the system so that IBM Spectrum Protect can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The IBM Spectrum Protect offering contains all of the media that IBM Installation Manager requires to install IBM Spectrum Protect.

Package

The group of software components that are required to install an offering.

The IBM Spectrum Protect package contains the following components:

- IBM Installation Manager installation program
- IBM Spectrum Protect offering

Package group

A set of packages that share a common parent directory.

The default package group for the IBM Spectrum Protect package is `IBM Installation Manager`.

Repository

A remote or local storage area for data and other application resources.

The IBM Spectrum Protect package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of IBM Spectrum Protect.

Windows: Worksheets for planning details for the server

You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect™ server. You can also use them to keep track of names and user IDs.

Windows Restriction: If you are using a File Allocation Table (FAT or FAT32) or a New Technology File System (NTFS) format, you cannot specify the root directory of that system as the location of a database directory or log directory. Instead, you must create one or more subdirectories within the root directory. Then, create the database directories and log directories within the subdirectories.

Item	Space required	Number of directories	Location of directories
The database			
Active log			
Archive log			
Optional: Log mirror for the active log			
Optional: Secondary archive log (failover location for archive log)			

Item	Names and user IDs	Location
------	--------------------	----------

Item	Names and user IDs	Location
The <i>instance user ID</i> for the server, which is the ID you use to start and run the IBM Spectrum Protect server		
The <i>home directory</i> for the server, which is the directory that contains the instance user ID		
The database instance name		
The <i>instance directory</i> for the server, which is a directory that contains files specifically for this server instance (the server options file and other server-specific files)		
The server name, use a unique name for each server		

Windows: Capacity planning

Capacity planning for IBM Spectrum Protect™ includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.

- **Windows: Estimating space requirements for the database**
To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.
- **Windows: Recovery log space requirements**
In IBM Spectrum Protect, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.
- **Windows: Monitoring space utilization for the database and recovery logs**
To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.
- **Windows: Deleting installation rollback files**
You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

Windows: Estimating space requirements for the database

To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.

About this task

Consider using at least 25 GB for the initial database space. Provision file system space appropriately. A database size of 25 GB is adequate for a test environment or a library-manager-only environment. For a production server supporting client workloads, the database size is expected to be larger. If you use random-access disk (DISK) storage pools, more database and log storage space is needed than for sequential-access storage pools.

The maximum size of the IBM Spectrum Protect™ database is 6 TB.

For information about sizing the database in a production environment that is based on the number of files and on storage pool size, see the following topics.

- **Windows: Estimating database space requirements based on the number of files**
If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

- Windows: Estimating database space requirements based on storage pool capacity
To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.
- Windows: The database manager and temporary space
The IBM Spectrum Protect server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

Windows: Estimating database space requirements based on the number of files

If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

About this task

To estimate space requirements that is based on the maximum number of files in server storage, use the following guidelines:

- 600 - 1000 bytes for each stored version of a file, including image backups.
Restriction: The guideline does not include space that is used during data deduplication.
- 100 - 200 bytes for each cached file, copy storage pool file, active-data pool file, and deduplicated file.
- Additional space is required for database optimization to support varying data-access patterns and to support server back-end processing of the data. The amount of extra space is equal to 50% of the estimate for the total number of bytes for file objects.

In the following example for a single client, the calculations are based on the maximum values in the preceding guidelines. The examples do not take into account that you might use file aggregation. In general, when you aggregate small files, it reduces the amount of required database space. File aggregation does not affect space-managed files.

Procedure

1. Calculate the number of file versions. Add each of the following values to obtain the number of file versions:
 - a. Calculate the number of backed-up files. For example, as many as 500,000 client files might be backed up at a time. In this example, storage policies are set to keep up to three copies of backed up files:

$$500,000 \text{ files} * 3 \text{ copies} = 1,500,000 \text{ files}$$

- b. Calculate the number of archive files. For example, as many as 100,000 client files might be archived copies.
- c. Calculate the number of space-managed files. For example, as many as 200,000 client files might be migrated from client workstations.

Using 1000 bytes per file, the total amount of database space that is required for the files that belong to the client is 1.8 GB:

$$(1,500,000 + 100,000 + 200,000) * 1000 = 1.8 \text{ GB}$$

2. Calculate the number of cached files, copy storage-pool files, active-data pool files, and deduplicated files:
 - a. Calculate the number of cached copies. For example, caching is enabled in a 5 GB disk storage pool. The high migration threshold of the pool is 90% and the low migration threshold of the pool is 70%. Thus, 20% of the disk pool, or 1 GB, is occupied by cached files.
If the average file size is about 10 KB, approximately 100,000 files are in cache at any one time:

$$100,000 \text{ files} * 200 \text{ bytes} = 19 \text{ MB}$$

- b. Calculate the number of copy storage-pool files. All primary storage pools are backed up to the copy storage pool:

$$(1,500,000 + 100,000 + 200,000) * 200 \text{ bytes} = 343 \text{ MB}$$

- c. Calculate the number of active storage-pool files. All the active client-backup data in primary storage pools is copied to the active-data storage pool. Assume that 500,000 versions of the 1,500,000 backup files in the primary storage pool are active:

$$500,000 * 200 \text{ bytes} = 95 \text{ MB}$$

d. Calculate the number of deduplicated files. Assume that a deduplicated storage pool contains 50,000 files:

$$50,000 * 200 \text{ bytes} = 10 \text{ MB}$$

Based on the preceding calculations, about 0.5 GB of extra database space is required for the client's cached files, copy storage-pool files, active-data pool files, and deduplicated files.

3. Calculate the amount of extra space that is required for database optimization. To provide optimal data access and management by the server, extra database space is required. The amount of extra database space is equal to 50% of the total space requirements for file objects.

$$(1.8 + 0.5) * 50\% = 1.2 \text{ GB}$$

4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

$$1.8 + 0.5 + 1.2 = 3.5 \text{ GB}$$

5. Calculate the total amount of database space that is required for all clients. If the client that was used in the preceding calculations is typical and you have 500 clients, for example, you can use the following calculation to estimate the total amount of database space that is required for all clients:

$$500 * 3.5 = 1.7 \text{ TB}$$

Results

Tip: In the preceding examples, the results are estimates. The actual size of the database might differ from the estimate because of factors such as the number of directories and the length of the path and file names. Periodically monitor your database and adjust its size as necessary.

What to do next

During normal operations, the IBM Spectrum Protect™ server might require temporary database space. This space is needed for the following reasons:

- To hold the results of sorting or ordering that are not already being kept and optimized in the database directly. The results are temporarily held in the database for processing.
- To give administrative access to the database through one of the following methods:
 - A DB2® open database connectivity (ODBC) client
 - An Oracle Java™ database connectivity (JDBC) client
 - Structured Query Language (SQL) to the server from an administrative-client command line

Consider using an extra 50 GB of temporary space for every 500 GB of space for file objects and optimization. See the guidelines in the following table. In the example that is used in the preceding step, a total of 1.7 TB of database space is required for file objects and optimization for 500 clients. Based on that calculation, 200 GB is required for temporary space. The total amount of required database space is 1.9 TB.

Database size	Minimum temporary-space requirement
< 500 GB	50 GB
≥ 500 GB and < 1 TB	100 GB
≥ 1 TB and < 1.5 TB	150 GB
≥ 1.5 and < 2 TB	200 GB
≥ 2 and < 3 TB	250 - 300 GB
≥ 3 and < 4 TB	350 - 400 GB

Windows: Estimating database space requirements based on storage pool capacity

To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.

Windows: The database manager and temporary space

The IBM Spectrum Protect™ server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

The database manager sorts data in a specific sequence, according to the SQL statement that you issue to request the data. Depending on the workload on the server, and if there is more data than the database manager can manage, the data (that is ordered in sequence) is allocated to temporary disk space. Data is allocated to temporary disk space when there is a large result set. The database manager dynamically manages the memory that is used when data is allocated to temporary disk space.

For example, expiration processing can produce a large result set. If there is not enough system memory on the database to store the result set, some of the data is allocated to temporary disk space. During expiration processing, if a node or file space are selected that are too large to process, the database manager cannot sort the data in memory. The database manager must use temporary space to sort data.

To run database operations, consider adding more database space for the following scenarios:

- The database has a small amount of space and the server operation that requires temporary space uses the remaining free space.
- The file spaces are large, or the file spaces have an assigned policy that creates many file versions.
- The IBM Spectrum Protect server must run with limited memory. The database uses the IBM Spectrum Protect server main memory to run database operations. However, if there is insufficient memory available, the IBM Spectrum Protect server allocates temporary space on disk to the database. For example, if 10G of memory is available and database operations require 12G of memory, the database uses temporary space.
- An `out of database space` error is displayed when you deploy an IBM Spectrum Protect server. Monitor the server activity log for messages that are related to database space.

Important: Do not change the DB2 software that is installed with the IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack, of DB2 software to avoid damage to the database.

Windows: Recovery log space requirements

In IBM Spectrum Protect™, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.

- **Windows: Active and archive log space**
When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.
- **Windows: Active-log mirror space**
The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.
- **Windows: Archive-failover log space**
The archive failover log is used by the server if the archive log directory runs out of space.

Windows: Active and archive log space

When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.

In IBM Spectrum Protect™ servers V7.1 and later, the active log can be a maximum size of 512 GB. The archive log size is limited to the size of the file system that it is installed on.

Use the following general guidelines when you estimate the size of the active log:

- The suggested starting size for the active log is 16 GB.
- Ensure that the active log is at least large enough for the amount of concurrent activity that the server typically handles. As a precaution, try to anticipate the largest amount of work that the server manages at one time. Provision the active log with extra space that can be used if needed. Consider using 20% of extra space.
- Monitor used and available active log space. Adjust the size of the active log as needed, depending upon factors such as client activity and the level of server operations.

- Ensure that the directory that holds the active log is as large as, or larger than, the size of the active log. A directory that is larger than the active log can accommodate failovers, if they occur.
- Ensure that the file system that contains the active log directory has at least 8 GB of free space for temporary log movement requirements.

The suggested starting size for the archive log is 48 GB.

The archive log directory must be large enough to contain the log files that are generated since the previous full backup. For example, if you perform a full backup of the database every day, the archive log directory must be large enough to hold the log files for all the client activity that occurs during 24 hours. To recover space, the server deletes obsolete archive log files after a full backup of the database. If the archive log directory becomes full and a directory for archive failover logs does not exist, log files remain in the active log directory. This condition can cause the active log directory to fill up and stop the server. When the server restarts, some of the existing active-log space is released.

After the server is installed, you can monitor archive log utilization and the space in the archive log directory. If the space in the archive log directory fills up, it can cause the following problems:

- The server is unable to perform full database backups. Investigate and resolve this problem.
- Other applications write to the archive log directory, exhausting the space that is required by the archive log. Do not share archive log space with other applications including other IBM Spectrum Protect servers. Ensure that each server has a separate storage location that is owned and managed by that specific server.
- Windows: Example: Estimating active and archive log sizes for basic client-store operations
Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.
- Windows: Example: Estimating active and archive log sizes for clients that use multiple sessions
If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.
- Windows: Example: Estimating active and archive log sizes for simultaneous write operations
If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.
- Windows: Example: Estimating active and archive log sizes for basic client store operations and server operations
Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.
- Windows: Example: Estimating active and archive log sizes under conditions of extreme variation
Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.
- Windows: Example: Estimating archive log sizes with full database backups
The IBM Spectrum Protect server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.
- Windows: Example: Estimating active and archive log sizes for data deduplication operations
If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

Windows: Example: Estimating active and archive log sizes for basic client-store operations

Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.

To determine the sizes of the active and archive logs for basic client-store operations, use the following calculation:

```
number of clients x files stored during each transaction
x log space needed for each file
```

This calculation is used in the example in the following table.

Table 1. Basic client-store operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053 bytes	The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	19.5 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 bytes = 3.5 GB Increase that amount by the suggested starting size of 16 GB: 3.5 + 16 = 19.5 GB
Archive log: Suggested size	58.5 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. 3.5 x 3 = 10.5 GB Increase that amount by the suggested starting size of 48 GB: 10.5 + 48 = 58.5 GB
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Windows: Example: Estimating active and archive log sizes for clients that use multiple sessions

If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.

To determine the sizes of the active and archive logs when clients use multiple sessions, use the following calculation:

number of clients x sessions for each client x files stored
during each transaction x log space needed for each file

This calculation is used in the example in the following table.

Table 1. Multiple client sessions

Item	Example values	Description
------	----------------	-------------

Item	Example values		Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	1000	The number of client nodes that back up, archive, or migrate files every night.
Possible sessions for each client	3	3	The setting of the client option RESOURCEUTILIZATION is larger than the default. Each client session runs a maximum of three sessions in parallel.
Files stored during each transaction	4096	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053	3053	<p>The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>
Active log: Suggested size	26.5 GB ¹	51 GB ¹	<p>The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 3 sessions for each client x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 10.5 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>10.5 + 16 = 26.5 GB</p> <p>The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.</p> <p>(1000 clients x 3 sessions for each client x 4096 files store during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 35 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>35 + 16 = 51 GB</p>
Archive log: Suggested size	79.5 GB ¹	153 GB ¹	<p>Because of the requirement to be able to store archive logs across three server-database backup cycles, the estimate for the active log is multiplied by 3:</p> <p>10.5 x 3 = 31.5 GB</p> <p>35 x 3 = 105 GB</p> <p>Increase those amounts by the suggested starting size of 48 GB:</p> <p>31.5 + 48 = 79.5 GB</p> <p>105 + 48 = 153 GB</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your active log and adjust its size if necessary.</p>			

Windows: Example: Estimating active and archive log sizes for simultaneous write operations

If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.

The log space that is required for each file increases by about 200 bytes for each copy storage pool that is used for a simultaneous write operation. In the example in the following table, data is stored to two copy storage pools in addition to a primary storage pool. The estimated log size increases by 400 bytes for each file. If you use the suggested value of 3053 bytes of log space for each file, the total number of required bytes is 3453.

This calculation is used in the example in the following table.

Table 1. Simultaneous write operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3453 bytes	3053 bytes plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files stored during each transaction x 3453 bytes for each file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size	60 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Windows: Example: Estimating active and archive log sizes for basic client store operations and server operations

Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.

For example, migration of files from the random-access (DISK) storage pool to a sequential-access disk (FILE) storage pool uses approximately 110 bytes of log space for each file that is migrated. For example, suppose that you have 300 backup-archive clients and each one of them backs up 100,000 files every night. The files are initially stored on DISK and then migrated to a FILE storage pool. To estimate the amount of active log space that is required for the data migration, use the following calculation. The number of clients in the calculation represents the maximum number of client nodes that back up, archive, or migrate files concurrently at any time.

```
300 clients x 100,000 files for each client x 110 bytes = 3.1 GB
```

Add this value to the estimate for the size of the active log that calculated for basic client store operations.

Windows: Example: Estimating active and archive log sizes under conditions of extreme variation

Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

Windows: Example: Estimating archive log sizes with full database backups

The IBM Spectrum Protect™ server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.

For example, if a full database backup occurs once a week, the archive log space must be able to contain the information in the archive log for a full week.

The difference in archive log size for daily and full database backups is shown in the example in the following table.

Table 1. Full database backups

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.

Item	Example values	Description
Log space that is required for each file	3453 bytes	3053 bytes for each file plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.
Active log: Suggested size	20 GB ¹	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes. (300 clients x 4096 files per transaction x 3453 bytes per file) ÷ 1,073,741,824 bytes = 4.0 GB Increase that amount by the suggested starting size of 16 GB: 4 + 16 = 20 GB
Archive log: Suggested size with a full database backup every day	60 GB ¹	Because of the requirement to be able to store archive logs across three backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement: 4 GB x 3 = 12 GB Increase that amount by the suggested starting size of 48 GB: 12 + 48 = 60 GB
Archive log: Suggested size with a full database every week	132 GB ¹	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. Multiply the result by the number of days between full database backups: (4 GB x 3) x 7 = 84 GB Increase that amount by the suggested starting size of 48 GB: 84 + 48 = 132 GB
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested starting size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Windows: Example: Estimating active and archive log sizes for data deduplication operations

If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

The following factors affect requirements for active and archive log space:

The amount of deduplicated data

The effect of data deduplication on the active log and archive log space depends on the percentage of data that is eligible for deduplication. If the percentage of data that can be deduplicated is relatively high, more log space is required.

The size and number of extents

Approximately 1,500 bytes of active log space are required for each extent that is identified by a duplicate-identification process. For example, if 250,000 extents are identified by a duplicate-identification process, the estimated size of the active log is 358 MB:

250,000 extents identified during each process x 1,500 bytes
for each extent = 358 MB

Consider the following scenario. Three hundred backup-archive clients back up 100,000 files each night. This activity creates a workload of 30,000,000 files. The average number of extents for each file is two. Therefore, the total number of extents is 60,000,000, and the space requirement for the archive log is 84 GB:

60,000,000 extents x 1,500 bytes for each extent = 84 GB

A duplicate-identification process operates on aggregates of files. An aggregate consists of files that are stored in a given transaction, as specified by the TXNGROUPMAX server option. Suppose that the TXNGROUPMAX server option is set to the default of 4096. If the average number of extents for each file is two, the total number of extents in each aggregate is 8192, and the space required for the active log is 12 MB:

8192 extents in each aggregate x 1500 bytes for each extent =
12 MB

The timing and number of the duplicate-identification processes

The timing and number of duplicate-identification processes also affects the size of the active log. Using the 12 MB active-log size that was calculated in the preceding example, the concurrent load on the active log is 120 MB if 10 duplicate-identification processes are running in parallel:

12 MB for each process x 10 processes = 120 MB

File size

Large files that are processed for duplicate identification can also affect the size of the active log. For example, suppose that a backup-archive client backs up an 80 GB, file-system image. This object can have a high number of duplicate extents if, for example, the files included in the file system image were backed up incrementally. For example, assume that a file system image has 1.2 million duplicate extents. The 1.2 million extents in this large file represent a single transaction for a duplicate-identification process. The total space in the active log that is required for this single object is 1.7 GB:

1,200,000 extents x 1,500 bytes for each extent = 1.7 GB

If other, smaller duplicate-identification processes occur at the same time as the duplicate-identification process for a single large object, the active log might not have enough space. For example, suppose that a storage pool is enabled for deduplication. The storage pool has a mixture of data, including many relatively small files that range from 10 KB to several hundred KB. The storage pool also has few large objects that have a high percentage of duplicate extents.

To take into account not only space requirements but also the timing and duration of concurrent transactions, increase the estimated size of the active log by a factor of two. For example, suppose that your calculations for space requirements are 25 GB (23.3 GB + 1.7 GB for deduplication of a large object). If deduplication processes are running concurrently, the suggested size of the active log is 50 GB. The suggested size of the archive log is 150 GB.

The examples in the following tables show calculations for active and archive logs. The example in the first table uses an average size of 700 KB for extents. The example in the second table uses an average size of 256 KB. As the examples show, the average deduplicate-extent size of 256 KB indicates a larger estimated size for the active log. To minimize or prevent operational problems for the server, use 256 KB to estimate the size of the active log in your production environment.

Table 1. Average duplicate-extent size of 700 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	700 KB	700 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average size for extents.

Item	Example values		Description
Extents for a given file	1,198,372 bits	6,135,667 bits	<p>Using the average extent size (700 KB), these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for an 800 GB object: $(800 \text{ GB} \div 700 \text{ KB}) = 1,198,372 \text{ bits}$</p> <p>The following calculation was used for a 4 TB object: $(4 \text{ TB} \div 700 \text{ KB}) = 6,135,667 \text{ bits}$</p>
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	1.7 GB	8.6 GB	<p>The estimated active log space that are needed for this transaction.</p>
Active log: Suggested total size	66 GB ¹	79.8 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of two. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object: $(23.3 \text{ GB} + 1.7 \text{ GB}) \times 2 = 50 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB: $50 + 16 = 66 \text{ GB}$</p> <p>The following calculation was used for multiple transactions and a 4 TB object: $(23.3 \text{ GB} + 8.6 \text{ GB}) \times 2 = 63.8 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB: $63.8 + 16 = 79.8 \text{ GB}$</p>

Item	Example values		Description
Archive log: Suggested size	198 GB ¹	239.4 GB ¹	<p>Multiply the estimated size of the active log by a factor of 3.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $50 \text{ GB} \times 3 = 150 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $150 + 48 = 198 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $63.8 \text{ GB} \times 3 = 191.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $191.4 + 48 = 239.4 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Table 2. Average duplicate-extent size of 256 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	256 KB	256 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average extent size.
Extents for a given file	3,276,800 bits	16,777,216 bits	<p>Using the average extent size, these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(800 \text{ GB} \div 256 \text{ KB}) = 3,276,800 \text{ bits}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(4 \text{ TB} \div 256 \text{ KB}) = 16,777,216 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	4.5 GB	23.4 GB	The estimated size of the active log space that is required for this transaction.

Item	Example values		Description
Active log: Suggested total size	71.6 GB ¹	109.4 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of 2. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 4.5 \text{ GB}) \times 2 = 55.6 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $55.6 + 16 = 71.6 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 23.4 \text{ GB}) \times 2 = 93.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $93.4 + 16 = 109.4 \text{ GB}$
Archive log: Suggested size	214.8 GB ¹	328.2 GB ¹	<p>The estimated size of the active log multiplied by a factor of 3.</p> <p>The following calculation was used for an 800 GB object:</p> $55.6 \text{ GB} \times 3 = 166.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $166.8 + 48 = 214.8 \text{ GB}$ <p>The following calculation was used for a 4 TB object:</p> $93.4 \text{ GB} \times 3 = 280.2 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $280.2 + 48 = 328.2 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Windows: Active-log mirror space

The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.

Creating a log mirror is a suggested option. If you increase the size of the active log, the log mirror size is increased automatically. Mirroring the log can affect performance because of the doubled I/O activity that is required to maintain the mirror. The additional space that the log mirror requires is another factor to consider when deciding whether to create a log mirror.

If the mirror log directory becomes full, the server issues error messages to the activity log and to the db2diag.log. Server activity continues.

Windows: Archive-failover log space

The archive failover log is used by the server if the archive log directory runs out of space.

Specifying an archive failover log directory can prevent problems that occur if the archive log runs out of space. If both the archive log directory and the drive or file system where the archive failover log directory is located become full, the data remains in the active log directory. This condition can cause the active log to fill up, which causes the server to halt.

Windows: Monitoring space utilization for the database and recovery logs

To determine the amount of used and available active log space, you issue the QUERY LOG command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

Active log

If the amount of available active log space is too low, the following messages are displayed in the activity log:

ANR4531I: IC_AUTOBACKUP_LOG_USED_SINCE_LAST_BACKUP_TRIGGER

This message is displayed when the active log space exceeds the maximum specified size. The IBM Spectrum Protect™ server starts a full database backup.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR0297I: IC_BACKUP_NEEDED_LOG_USED_SINCE_LAST_BACKUP

This message is displayed when the active log space exceeds the maximum specified size. You must back up the database manually.

To change the maximum log size, halt the server. Open the dsmserv.opt file, and specify a new value for the ACTIVELOGSIZE option. When you are finished, restart the server.

ANR4529I: IC_AUTOBACKUP_LOG_UTILIZATION_TRIGGER

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. If at least one full database backup has occurred, the IBM Spectrum Protect server starts an incremental database backup. Otherwise, the server starts a full database backup.

ANR0295I: IC_BACKUP_NEEDED_LOG_UTILIZATION

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. You must back up the database manually.

Archive log

If the amount of available archive log space is too low, the following message is displayed in the activity log:

ANR0299I: IC_BACKUP_NEEDED_ARCHLOG_USED

The ratio of used archive-log space to available archive-log space exceeds the log utilization threshold. The IBM Spectrum Protect server starts a full automatic database backup.

Database

If the amount of space available for database activities is too low, the following messages are displayed in the activity log:

ANR2992W: IC_LOG_FILE_SYSTEM_UTILIZATION_WARNING_2

The used database space exceeds the threshold for database space utilization. To increase the space for the database, use the EXTEND DBSPACE command, the EXTEND DBSPACE command, or the DSMSERV FORMAT utility with the DBDIR parameter.

ANR1546W: FILESYSTEM_DBPATH_LESS_1GB

The available space in the directory where the server database files are located is less than 1 GB.

When an IBM Spectrum Protect server is created with the DSMSEV FORMAT utility or with the configuration wizard, a server database and recovery log are also created. In addition, files are created to hold database information used by the database manager. The path specified in this message indicates the location of the database information used by the database manager. If space is unavailable in the path, the server can no longer function.

You must add space to the file system or make space available on the file system or disk.

Windows: Deleting installation rollback files

You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

About this task

To delete the files that are no longer needed, use either the installation graphical wizard or the command line in console mode.

- **Windows: Deleting installation rollback files by using a graphical wizard**
You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.
- **Windows: Deleting installation rollback files by using the command line**
You can delete certain installation files that were saved during the installation process by using the command line.

Windows: Deleting installation rollback files by using a graphical wizard

You can delete certain installation files that were saved during installation process by using the IBM® Installation Manager user interface.

Procedure

1. Open IBM Installation Manager.
2. Click File > Preferences.
3. Select Files for Rollback.
4. Click Delete Saved Files and click OK.

Windows: Deleting installation rollback files by using the command line

You can delete certain installation files that were saved during the installation process by using the command line.

Procedure

1. In the directory where IBM® Installation Manager is installed, go to the following subdirectory:
 - o **Windows** eclipse\tools

For example:

- o **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools
2. From the tools directory, issue the following command to start an IBM Installation Manager command line:
 - o **Windows** imcl.exe -c
 3. Enter **P** to select Preferences.
 4. Enter **3** to select Files for Rollback.
 5. Enter **D** to Delete the Files for Rollback.
 6. Enter **A** to Apply Changes and Return to Preferences Menu.
 7. Enter **C** to leave the Preference Menu.
 8. Enter **X** to Exit Installation Manager.

Windows: Server naming best practices

Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect™ server.

Instance user ID

The instance user ID is used as the basis for other names related to the server instance. The instance user ID is also called the instance owner.

For example: tsminst1

The instance user ID is the user ID that must have ownership or read/write access authority to all directories that you create for the database and the recovery log. The standard way to run the server is under the instance user ID. That user ID must also have read/write access to the directories that are used for any FILE device classes.

Windows

Database instance name

The database instance name is the name of the server instance as it appears in the registry.

For example: Server1

Windows

Instance directory

The instance directory is a directory that contains files specifically for a server instance (the server options file and other server-specific files). It can have any name that you want. For easier identification, use a name that ties the directory to the instance name.

You can use a name that includes the name of the server instance as it appears (or will appear) in the registry. Default server instance names have the form Serverx.

For example: C:\tsm\server1

The instance directory stores the following files for the server instance:

- The server options file, dsmserv.opt
- The server key database file, cert.kdb, and the .arm files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the DEVCONFIG server option does not specify a fully qualified name
- Volume history file, if the VOLUMEHISTORY server option does not specify a fully qualified name
- Volumes for DEVTYPE=FILE storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)

Database name

The database name is always TSMDB1, for every server instance. This name cannot be changed.

Server name

The server name is an internal name for IBM Spectrum Protect, and is used for operations that involve communication among multiple IBM Spectrum Protect servers. Examples include server-to-server communication and library sharing.

Windows The server name is also used when you add the server to the Operations Center so that it can be managed using that interface. Use a unique name for each server. For easy identification in the Operations Center (or from a QUERY SERVER command), use a name that reflects the location or purpose of the server. Do not change the name of an IBM Spectrum Protect server after it is configured as a hub or spoke server.

If you use the wizard, the default name that is suggested is the host name of the system that you are using. You can use a different name that is meaningful in your environment. If you have more than one server on the system and you use the wizard, you can use the default name for only one of the servers. You must enter a unique name for each server.

Windows

For example,

- TUCSON_SERVER1

- TUCSON_SERVER2

Directories for database space and recovery log

The directories can be named according to local practices. For easier identification, consider using names that tie the directories to the server instance.

For example, for the archive log:

- **Windows** f:\server1\archlog

Windows: Installation directories

Installation directories for the IBM Spectrum Protect™ server include the server, DB2®, device, language, and other directories. Each one contains several additional directories.

The (/opt/tivoli/tsm/server/bin) is the default directory that contains server code and licensing.

The DB2 product that is installed as part of the installation of the IBM Spectrum Protect server has the directory structure as documented in DB2 information sources. Protect these directories and files as you do the server directories. The default directory is /opt/tivoli/tsm/db2.

You can use US English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, and Russian.

Windows: Installing the server components

To install the Version 8.1.5 server components, you can use the installation wizard, the command line in console mode, or silent mode.

About this task

Using the IBM Spectrum Protect™ installation software, you can install the following components:

- server
Tip: The database (DB2®), the Global Security Kit (GSKit) and IBM® Java™ Runtime Environment (JRE) are automatically installed when you select the server component.
- server languages
- license
- devices
- IBM Spectrum Protect for SAN
- Operations Center

Windows Allow approximately 15 - 30 minutes to install a V8.1.5 server, using this guide.

- Windows: Obtaining the installation package
You can obtain the IBM Spectrum Protect installation package from an IBM download site such as Passport Advantage® or IBM Fix Central.
- Windows: Installing IBM Spectrum Protect by using the installation wizard
You can install the server by using the IBM Installation Manager graphical wizard.
- Windows: Installing IBM Spectrum Protect by using console mode
You can install IBM Spectrum Protect by using the command line in console mode.
- Windows: Installing IBM Spectrum Protect in silent mode
You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.
- Windows: Installing server language packages
Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Windows: Obtaining the installation package

You can obtain the IBM Spectrum Protect™ installation package from an IBM® download site such as Passport Advantage® or IBM Fix Central.

Procedure

1. Download the appropriate package file from one of the following websites.
 - o Download the server package from Passport Advantage or Fix Central.
 - o For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you downloaded the package from an IBM download site, complete the following steps:
 - Windows**
 - a. Verify that you have enough space to store the installation files when they are extracted from the product package. See the download document for the space requirements:
 - IBM Spectrum Protect technote 4042944
 - IBM Spectrum Protect Extended Edition technote 4042945
 - IBM Spectrum Protect for Data Retention technote 4042946
 - b. Change to the directory where you placed the executable file.

Important: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.
 - c. Either double-click the executable file, or enter the following command on the command line to extract the installation files. The files are extracted to the current directory.

package_name.exe

where *package_name* is like this example: 8.1.x.000-IBM-SPSRV-WindowsX64.exe
3. Select one of the following methods of installing IBM Spectrum Protect:
 - o Windows: Installing IBM Spectrum Protect by using the installation wizard
 - o Windows: Installing IBM Spectrum Protect by using console mode
 - o Windows: Installing IBM Spectrum Protect in silent mode
4. After you install IBM Spectrum Protect, and before you customize it for your use, go to the IBM Support Portal. Click Support and downloads and apply any applicable fixes.

Windows: Installing IBM Spectrum Protect by using the installation wizard

You can install the server by using the IBM® Installation Manager graphical wizard.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- **Windows** Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

Install IBM Spectrum Protect™ by using this method:

Option	Description
Installing the software from a downloaded package:	<ol style="list-style-type: none">a. Change to the directory where you downloaded the package.b. Start the installation wizard by issuing the following command: Windows <code>install.bat</code> Windows Or, in the directory where the installation files were extracted, double-click the install.bat file.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

You can view installation log files by clicking File > View Log from the Installation Manager tool. To collect these log files, click Help > Export Data for Problem Analysis from the Installation Manager tool.

- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **Windows** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.
- **Windows** If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Spectrum Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is `C:\Program Files\Tivoli\TSM\device\drivers`.

Windows: Installing IBM Spectrum Protect by using console mode

You can install IBM Spectrum Protect™ by using the command line in console mode.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- **Windows** Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

Install IBM Spectrum Protect by using this method:

Option	Description
Installing the software from a downloaded package:	<p>a. Change to the directory where you downloaded the package.</p> <p>b. Start the installation wizard in console mode by issuing the following command: Windows</p> <pre>install.bat -c</pre> <p>Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the Summary panel, specify <code>G</code> to generate the responses.</p>

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - **Windows** `C:\ProgramData\IBM\Installation Manager\logs`
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **Windows** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.
- **Windows** If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Spectrum Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is `C:\Program Files\Tivoli\TSM\device\drivers`.

Windows: Installing IBM Spectrum Protect in silent mode

You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the input directory where the installation package is extracted:

install_response_sample.xml

Use this file to install the IBM Spectrum Protect™ components.

update_response_sample.xml

Use this file to upgrade the IBM Spectrum Protect components.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
2. If you install the server and Operations Center in silent mode, create a password for the Operations Center truststore in the response file.
If you are using the install_response_sample.xml file, add the password in the following line of the file, where *mypassword* represents the password:

```
<variable name='ssl.password' value='mypassword' />
```

For more information about this password, see Installation checklist

Tip: To upgrade the Operations Center, the truststore password is not required if you are using the update_response_sample.xml file.

3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:

o **Windows**

```
install.bat -s -input response_file -acceptLicense
```

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM® Installation Manager logs directory, for example:
 - o **Windows** C:\ProgramData\IBM\Installation Manager\logs
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click Downloads (fixes and PTFs) and apply any applicable fixes.
- **Windows** After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.
- **Windows** If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Spectrum Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is C:\Program Files\Tivoli\TSM\device\drivers.

Windows

Windows: Installing server language packages

Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Before you begin

For instructions on installing storage agent language packages, see Language pack configuration for storage agents.

- Windows: Server language locales
Use either the default language package option or select another language package to display server messages and help.

- Windows: Configuring a language package
After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.
- Windows: Updating a language package
You can modify or update a language package by using the IBM® Installation Manager.

Windows: Server language locales

Use either the default language package option or select another language package to display server messages and help.

Windows This language package is automatically installed for the following default language option for server messages and help: LANGUAGE AMENG.

For languages or locales other than the default, install the language package that your installation requires. You can use the languages that are shown:

Windows
Table 1. Server languages for Windows

Language	LANGUAGE option value
Chinese, Simplified	chs
Chinese, Traditional	cht
English	ameng
French	fra
German	deu
Italian	ita
Japanese (Shift-JIS)	jpn
Korean	kor
Portuguese, Brazilian	ptb
Russian	rus
Spanish	esp

Windows Restriction: For Operations Center users, some characters might not be displayed properly if the web browser does not use the same language as the server. If this problem occurs, set the browser to use the same language as the server.

Windows: Configuring a language package

After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect™.

About this task

Windows Set the LANGUAGE option in the server options file to the name of the locale that you want to use. For example: to use the `ita` locale, set the LANGUAGE option to `ita`. See Windows: Server language locales.

If the locale is successfully initialized, it formats the date, time, and number for the server. If the locale is not successfully initialized, the server uses the US English message files and the date, time, and number format.

Windows: Updating a language package

You can modify or update a language package by using the IBM® Installation Manager.

About this task

You can install another language package within the same IBM Spectrum Protect™ instance.

- Use the Modify function of IBM Installation Manager to install another language package.

- Use the Update function of IBM Installation Manager to update to newer versions of the language packages.

Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.

Windows: Taking the first steps after you install IBM Spectrum Protect

After you install Version 8.1.5, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect™ instance.

About this task

1. Create the directories and user ID for the server instance. See [Windows: Creating the user ID and directories for the server instance](#).
2. Configure a server instance. Select one of the following options:
 - Use the configuration wizard, the preferred method. See [Windows: Configuring IBM Spectrum Protect by using the configuration wizard](#).
 - Manually configure the new instance. See [Windows: Configuring the server instance manually](#). Complete the following steps during a manual configuration.
 - a. Set up your directories and create the IBM Spectrum Protect instance. See [Windows: Creating the server instance](#).
 - b. Create a new server options file by copying the sample file to set up communications between the server and clients. See [Windows: Configuring server and client communications](#).
 - c. Issue the DSMSEV FORMAT command to format the database. See [Windows: Formatting the database and log](#).
 - d. Configure your system for database backup. See [Windows: Preparing the database manager for database backup](#).
3. Configure options to control when database reorganization runs. See [Windows: Configuring server options for server database maintenance](#).
4. Start the server instance if it is not already started.
 - [Windows](#) See [Windows: Starting the server instance on Windows systems](#).
5. Register your license. See [Windows: Registering licenses](#).
6. Prepare your system for database backups. See [Windows: Preparing the server for database backup operations](#).
7. Monitor the server. See [Windows: Monitoring the server](#).

- [Windows: Creating the user ID and directories for the server instance](#)
Create the user ID for the IBM Spectrum Protect server instance and create the directories that the server instance needs for database and recovery logs.
- [Windows: Configuring the IBM Spectrum Protect server](#)
After you have installed the server and prepared for the configuration, configure the server instance.
- [Windows: Configuring server options for server database maintenance](#)
To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.
- [Windows](#) [Windows: Starting the server instance on Windows systems](#)
In a production environment, the preferred method for starting the server is as a Windows service. In an environment where you are reconfiguring, testing, or completing maintenance tasks, start the server in the foreground or use maintenance mode.
- [Windows: Stopping the server](#)
You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.
- [Windows: Registering licenses](#)
Immediately register any IBM Spectrum Protect licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.
- [Windows: Preparing the server for database backup operations](#)
To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.
- [Windows: Running multiple server instances on a single system](#)
You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

- Windows: Monitoring the server
When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Windows: Creating the user ID and directories for the server instance

Create the user ID for the IBM Spectrum Protect™ server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See Windows: Worksheets for planning details for the server.

Procedure

1. Create the user ID that will own the server instance. You use this user ID when you create the server instance in a later step.

Windows

Windows Create a user ID that will be the owner of the IBM Spectrum Protect server instance. A user ID can own more than one IBM Spectrum Protect server instance. Identify the user account that will own the server instance.

When the server is started as a Windows service, this account is the one that the service will log on to. The user account must have administrative authority on the system. One user account can own more than one server instance.

If you have multiple servers on one system and want to run each server with a different user account, create a new user account in this step.

Create the user ID.

Restriction: The user ID must comply with the following rule:

In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID must be 30 characters or less, and cannot start with *ibm*, *sql*, *sys*, or a numeral. The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

- a. Use the following operating system command to create the user ID:

```
net user user_ID * /add
```

You are prompted to create and verify a password for the new user ID.

- b. Issue the following operating system commands to add the new user ID to the Administrators groups:

```
net localgroup Administrators user_ID /add
net localgroup DB2ADMNS user_ID /add
```

2. Create directories that the server requires.

Windows Create empty directories for each item in the table and ensure that the new user ID you just created has read/write permission to the directories. The database, archive log, and active log must reside on different physical volumes.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	<code>mkdir d:\tsm\server1</code>	
The database directories	<code>mkdir d:\tsm\db001</code> <code>mkdir e:\tsm\db002</code> <code>mkdir f:\tsm\db003</code> <code>mkdir g:\tsm\db004</code>	

Item	Example commands for creating the directories	Your directories
Active log directory	<code>mkdir h:\tsm\log</code>	
Archive log directory	<code>mkdir i:\tsm\archlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir j:\tsm\logmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir k:\tsm\archlogfailover</code>	

When a server is initially created by using the DSMSEV FORMAT utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

Windows: Configuring the IBM Spectrum Protect server

After you have installed the server and prepared for the configuration, configure the server instance.

About this task

Configure an IBM Spectrum Protect™ server instance by selecting one of the following options:

- **Windows: Configuring IBM Spectrum Protect by using the configuration wizard**
The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect server program.
- **Windows: Configuring the server instance manually**
After installing IBM Spectrum Protect Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

Windows: Configuring IBM Spectrum Protect by using the configuration wizard

The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect™ server program.

Before you begin

Before you use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing IBM Spectrum Protect, creating the database and log directories, and creating the directories and user ID for the server instance.

Windows

About this task

Procedure

1. Ensure that the following requirements are met:
 - A backup copy of the following files must be saved to a safe and secure location:
 - Master encryption key files (`dsmkeydb.*`)
 - Server certificate and private key files (`cert.*`)
 - Start the Remote Registry service:
 - a. Click Start > Administrative Tools > Services.
 - b. In the Services window, select the Remote Registry service if it is not started, and click Start.

Windows

- o Ensure that the following ports are not blocked by a firewall: 137, 139 and 445. Complete the following steps:
 - a. Click Start > Control Panel > Windows Firewall.
 - b. Select Advanced Settings.
 - c. Select Inbound Rules in the left pane.
 - d. Select New Rule in the right pane.
 - e. Create a port rule for TCP ports 137, 139 and 445 to allow connections for domain and private networks.
- o Configure User Account Control:

Access all three of the user account control configuration settings by first accessing Local Security Policy Security options. Complete the following steps:

- a. Enable the built-in Administrator account:
 - Select the Accounts: Administrator account status.
 - Select Enable and click OK.
 - b. Disable User Account Control for all Windows administrators:
 - Select the User Account Control: Run all administrators in Admin Approval Mode.
 - Select Disable and click OK.
 - c. Disable User Account Control for the built-in Administrator account:
 - Select the User Account Control: Admin Approval Mode for the Built-in Administrator Account.
 - Select Disable and click OK.
2. Start the local version of the wizard:
- o **Windows** Either click Start > All Programs > IBM Spectrum Protect > Configuration Wizard. Or, double-click the `dsmicfgx.exe` program in `installation_directory\server`. The default directory is `C:\Program Files\Tivoli\TSM`.

Follow the instructions to complete the configuration. The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

- **Windows** Windows: Configuring Remote Execution Protocol on Windows
Configure remote access settings by using these procedures.

Windows: Configuring the server instance manually

After installing IBM Spectrum Protect™ Version 8.1.5, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

- Windows: Creating the server instance
Create an IBM Spectrum Protect instance by issuing the `db2icrt` command.
- **Windows** Windows: Configuring server and client communications
After installing the server, you can set up client and server communications by specifying options in the server and client options files.
- Windows: Formatting the database and log
Use the `DSMSERV FORMAT` utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.
- Windows: Preparing the database manager for database backup
To back up the data in the database to IBM Spectrum Protect, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

Windows: Creating the server instance

Create an IBM Spectrum Protect™ instance by issuing the `db2icrt` command.

About this task

You can have one or more server instances on one workstation.

Windows Important: Before you run the `db2icrt` command, verify the following items:

- Ensure that the user and the instance directory of the user exists. If there is no instance directory, you must create it. The instance directory stores the following files for the server instance:
 - o The server options file, `dsmerv.opt`
 - o The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
 - o Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name

- o Volume history file, if the VOLUMEHISTORY server option does not specify a fully qualified name
- o Volumes for DEVTYPE=FILE storage pools, if the directory for the device class is not fully specified, or not fully qualified
- o User exits
- o Trace output (if not fully qualified)
- Save a backup copy of the following files to a safe and secure location:
 - o Master encryption key files (`dsmkeydb.*`)
 - o Server certificate and private key files (`cert.*`)

Windows

1. Log in as an administrator and create an IBM Spectrum Protect instance, by using the `db2icrt` command. Enter the following command on one line. The user account that you specify becomes the user ID that owns the Version 8.1.5 server (the instance user ID).

```
db2icrt -u user_account instance_name
```

For example, if the user account is `tsminst1` and the server instance is `Server1`, enter the following command:

```
db2icrt -u tsminst1 server1
```

You are prompted for the password for user ID `tsminst1`. Later, when you create and format the database, you use the instance name that you specified with this command, with the `-k` option.

2. Change the default path for the database to be the drive where the instance directory for the server is located. Complete the following steps:
 - a. Click Start > Programs > IBM DB2 > DB2TSM1 > Command Line Tools > Command Line Processor.
 - b. Enter `quit` to exit the command line processor.

A window with a command prompt should now be open, with the environment properly set up to successfully issue the commands in the next steps.

- c. From the command prompt in that window, issue the following command to set the environment variable for the server instance that you are working with:

```
set db2instance=instance_name
```

The `instance_name` is the same as the instance name that you specified when you issued the `db2icrt` command. For example, to set the environment variable for the `Server1` server instance, issue the following command:

```
set db2instance=server1
```

- d. Issue the command to set the default drive:

```
db2 update dbm cfg using dftdbpath instance_location
```

For example, the instance directory is `d:\tsm\server1` and the instance location is drive `d:`. Enter the command:

```
db2 update dbm cfg using dftdbpath d:
```

3. Create a new server options file. See Windows: Configuring server and client communications.

Windows

Windows: Configuring server and client communications

After installing the server, you can set up client and server communications by specifying options in the server and client options files.

About this task

Set these server options before you start the server. When you start the server, the new options go into effect. If you modify any server options after starting the server, you must stop and restart the server to activate the updated options.

Review the server options file (`dsmserve.opt.smp`) that is located in the server instance directory to view and specify server communications options. By default, the server uses the TCP/IP and Named Pipes communication methods.

Tip: If you start the server console and see warning messages that a protocol could not be used by the server, either the protocol is not installed or the settings do not match the Windows protocol settings.

For a client to use a protocol that is enabled on the server, the client options file must contain corresponding values for communication options. In the server options file, you can view the values for each protocol.

You can specify one or more of the following communication methods:

- TCP/IP Version 4 or Version 6
- Named Pipes
- Shared memory
- Secure Sockets Layer (SSL)

Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

- **Windows** Windows: Setting TCP/IP options
Select from a range of TCP/IP options for the IBM Spectrum Protect server or retain the default.
- **Windows** Windows: Setting Named Pipes options
The Named Pipes communication method is ideal when running the server and client on the same Windows machine. Named Pipes require no special configuration.
- **Windows** Windows: Setting Secure Sockets Layer options
You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Windows: Setting TCP/IP options

Select from a range of TCP/IP options for the IBM Spectrum Protect™ server or retain the default.

About this task

The following is an example of a list of TCP/IP options that you can use to set up your system.

```
commmethod      tcpip
tcpport         1500
tcpwindowsize   0
tcpnodelay      yes
```

Tip: You can use TCP/IP Version 4, Version 6, or both.

TCPPORT

The server port address for TCP/IP and SSL communication. The default value is 1500.

Windows TCPWINDOWSIZE

Windows Specifies the size of the TCP/IP buffer that is used when sending or receiving data. The window size that is used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but can improve performance.

To use the default window size for the operating system, specify 0.

TCPNODELAY

Specifies whether or not the server sends small messages or lets TCP/IP buffer the messages. Sending small messages can improve throughput but increases the number of packets sent over the network. Specify YES to send small messages or NO to let TCP/IP buffer them. The default is YES.

TCPADMINPORT

Specifies the port number on which the server TCP/IP communication driver is to wait for TCP/IP or SSL-enabled communication requests other than client sessions. The default is the value of TCPPORT.

SSLTCPPOINT

(SSL-only) Specifies the Secure Sockets Layer (SSL) port number on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line backup-archive client and the command-line administrative client.

SSLTCPADMINPORT

(SSL-only) Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client.

Windows: Setting Named Pipes options

The Named Pipes communication method is ideal when running the server and client on the same Windows machine. Named Pipes require no special configuration.

About this task

Here is an example of a Named Pipes setting:

```
commmethod          namedpipe
namedpipename       \\.\pipe\adsmpipe
```

COMMETHOD can be used multiple times in the IBM Spectrum Protect™ server options file, with a different value each time. For example, the following example is possible:

```
commmethod tcpip
commmethod namedpipe
```

Windows: Setting Secure Sockets Layer options

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Before you begin

SSL is the standard technology for creating encrypted sessions between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communication paths. With SSL, the identity of the server is verified through the use of digital certificates.

To ensure better system performance, use SSL only for sessions when it is needed. Consider adding additional processor resources on the IBM Spectrum Protect™ server to manage the increased requirements.

Windows: Formatting the database and log

Use the DSMSEV FORMAT utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

After you set up server communications, you are ready to initialize the database. Ensure that you log in by using the instance user ID. Do not place the directories on file systems that might run out of space. If certain directories (for example, the archive log) become unavailable or full, the server stops.

Windows Important: The installation program creates a set of registry keys. One of these keys points to the directory where a default server, named SERVER1, is created. To install an additional server, create a directory and use the DSMSEV FORMAT utility, with the -k parameter, from that directory. That directory becomes the location of the server. The registry tracks the installed servers.

Setting the exit list handler

Set the DB2NOEXITLIST registry variable to ON for each server instance. Log on to the system as the server instance owner and issue this command:

```
db2set -i server_instance_name DB2NOEXITLIST=ON
```

For example: **Windows**

```
db2set -i server1 DB2NOEXITLIST=ON
```

Initializing a server instance

Use the DSMSEV FORMAT utility to initialize a server instance. For example, if the server instance directory is */tsminst1*, issue the following commands: **Windows**

```
cd \tsminst1
dsmserve -k server2 format dbdir=d:\tsm\db001 activelogsiz=32768
activelogdirectory=e:\tsm\activelog archlogdirectory=f:\tsm\archlog
archfailoverlogdirectory=g:\tsm\archfaillog mirrorlogdirectory=h:\tsm\mirrorlog
```

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

Related information:

[DSMSERV FORMAT](#) (Format the database and log)

Windows: Preparing the database manager for database backup

To back up the data in the database to IBM Spectrum Protect™, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

About this task

If you use the configuration wizard to create an IBM Spectrum Protect server instance, you do not have to complete these steps. If you are configuring an instance manually, complete the following steps before you issue either the BACKUP DB or the RESTORE DB commands.

Attention: If the database is unusable, the entire IBM Spectrum Protect server is unavailable. If a database is lost and cannot be recovered, it might be difficult or impossible to recover data that is managed by that server. Therefore, it is critically important to back up the database.

Windows Restriction: Database backup and restore over shared memory are not available on Windows systems.

Windows In the following commands, the examples use `server1` for the database instance and `d:\tmsserver1` for the IBM Spectrum Protect server directory. Replace these values with your actual values in the commands.

1. Create a file that is called `tsmdbmgr.env` in the `d:\tmsserver1` directory with the following contents:

```
DSMI_CONFIG=server_instance_directory\tsmdbmgr.opt
DSMI_LOG=server_instance_directory
```

2. Set the `DSMI_api` environment-variable configuration for the database instance:

- a. Open a DB2® command window. One method is to go to the `C:\Program Files\Tivoli\TSM\db2\bin` directory, or if you installed IBM Spectrum Protect in a different location, go to the `db2\bin` subdirectory in your main installation directory. Then, issue this command:

```
db2cmd
```

- b. Issue this command:

```
db2set -i server1 DB2_VENDOR_INI=d:\tmsserver1\tsmdbmgr.env
```

3. Create a file that is called `tsmdbmgr.opt` in the `d:\tmsserver1` directory with the following contents:

```
*****
nodename $$_TSMDBMGR_$$
commethod tcpip
tcpserveraddr localhost
tcpport 1500
passwordaccess generate
errorlogname d:\tmsserver1\tsmdbmgr.log
```

where

- o `nodename` specifies the node name the client API uses to connect to the server during a database backup. This value must be `$$_TSMDBMGR_$$` for database backup to work.
- o `commethod` specifies the client API used to contact the server for database backup.
- o `tcpserveraddr` specifies the server address that the client API uses to contact the server for database backup. To ensure that the database can be backed up, this value must be `localhost`.
- o `tcpport` specifies the port number that the client API uses to contact the server for database backup. Ensure that you enter the same `tcpport` value that is specified in the `dsmserv.opt` server options file.
- o `passwordaccess` is required for the backup node to connect to the server on windows systems.
- o `errorlogname` specifies the error log where the client API logs errors that are encountered during a database backup. This log is typically in the server instance directory. However, this log can be placed in any location where the instance user ID has write-permission.

Windows: Configuring server options for server database maintenance

To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.

About this task

Table and index reorganization requires significant processor resources, active log space, and archive log space. Because database backup takes precedence over reorganization, select the time and duration for reorganization to ensure that the processes do not overlap and reorganization can complete.

Windows You can optimize index and table reorganization for the server database. In this way, you can help to avoid unexpected database growth and performance issues. For instructions, see technote 1683633.

If you update these server options while the server is running, you must stop and restart the server before the updated values take effect.

Procedure

1. Modify the server options.

Windows Edit the server options file, `dsmserv.opt`, in the server instance directory by using a text editor. Follow these guidelines when you edit the server options file:

- o To enable an option, remove the asterisk at the beginning of the line.
- o Enter an option on any line.
- o Enter only one option per line. The entire option with its value must be on one line.
- o If you have multiple entries for an option in the file, the server uses the last entry.

To view available server options, see the sample file, `dsmserv.opt.smp`, in the `c:\Program Files\Tivoli\TSM` directory.

2. If you plan to use data deduplication, enable the `ALLOWREORGINDEX` server option. Add the following option and value to the server options file:

```
allowreorgindex yes
```

3. Set the `REORGBEGINTIME` and `REORGDURATION` server options to control when reorganization starts and how long it runs. Select a time and duration so that reorganization runs when you expect that the server is least busy. These server options control both table and index reorganization processes.

- a. Set the time for reorganization to start by using the `REORGBEGINTIME` server option. Specify the time by using the 24-hour system. For example, to set the start time for reorganization as 8:30 p.m., specify the following option and value in the server options file:

```
reorgbegintime 20:30
```

- b. Set the interval during which the server can start reorganization. For example, to specify that the server can start reorganization for four hours after the time set by the `REORGBEGINTIME` server option, specify the following option and value in the server options file:

```
reorgduration 4
```

4. If the server was running while you updated the server options file, stop and restart the server.

Windows

Windows: Starting the server instance on Windows systems

In a production environment, the preferred method for starting the server is as a Windows service. In an environment where you are reconfiguring, testing, or completing maintenance tasks, start the server in the foreground or use maintenance mode.

Before you begin

Select one of the following methods for starting the server:

As a Windows service

This method is useful in a production environment. When you configure the server to run as a service, you can specify that the server starts automatically whenever the system is started.

In the foreground

This method is useful when you are configuring or testing the server. When you start the server in the foreground, IBM Spectrum Protect™ provides a special administrator user ID that is named SERVER_CONSOLE. All server messages are displayed in the foreground. The messages can be useful if you must debug startup problems.

In maintenance mode

This method is useful when you are completing maintenance or reconfiguration tasks. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Procedure

Follow the instructions for your selected option:

Option	Description
Starting the server as a Windows service	To start the server as a Windows service, take one of the following actions: <ul style="list-style-type: none"> • If you configured the server by using the configuration wizard, complete the following steps: <ol style="list-style-type: none"> a. Configure the server to start as a Windows service by following the instructions in Windows: Configuring the server to start as a Windows service. b. Start the server by following the instructions in Windows: Starting the server as a Windows service. • If you did not use the configuration wizard, create and configure the Windows service by following the instructions in Windows: Manually creating and configuring a Windows service.
Starting the server in the foreground	To start the server in the foreground, follow the instructions in Windows: Starting the server in the foreground.
Starting the server in maintenance mode	To start the server in maintenance mode, follow the instructions in Windows: Starting the server in maintenance mode.

Windows

Windows: Configuring the server to start as a Windows service

Before you can start the server as a Windows service, you must ensure that options and access rights are set correctly.

Before you begin

A Windows service must be created. If you configured the server by using the configuration wizard, a Windows service was created automatically. In that case, use this procedure to configure the server to start as a Windows service.

If you did not use a wizard, you must create and configure the Windows service manually by following the steps in Windows: Manually creating and configuring a Windows service.

Procedure

1. From the Windows Start menu, click Run, type `services.msc`, and click OK.
2. In the Services window, select the server instance that you want to start as a service, and click Properties. For example, select TSM INST1, and click Properties.
3. To ensure that the server service starts automatically, click the General tab. From the Startup type list, select Automatic.
4. To set the user for starting the server service, click the Log On tab, and take one of the following actions:
 - If you plan to run the server service under the Local System account, select Local System account and click OK.
 - If you plan to run the server service under the instance user ID, take the following actions:
 - a. Select This account, and browse for the user ID that owns the server DB2® instance and has permissions for starting the server.
 - b. In the Select User window, in the Enter the object name to select field, enter the user ID.
 - c. Click Check Names.
 - d. Click OK twice.
5. If you configured the server service to run under the Local System account, grant database access to the Local System account:
 - a. Log on with the user ID that was used to create the server database. This user ID is the user ID that was used to run the DSMSEV FORMAT utility to initialize the server database. Alternatively, if you configured the server with the dsmicfgx configuration wizard, this user ID is the user ID that was used to create the instance.
 - b. Open a DB2 command window. If the server is installed on Windows Server 2012, open the Start window, and click DB2 Command Window - Administrator.

c. In the DB2 command window, enter the following commands:

```
set DB2INSTANCE=server1
db2 connect to TSMDB1
db2 grant dbadm with dataaccess with accessctrl on database to user system
db2 grant secadm on database to user system
```

Tip: When the server service is configured to run under the Local System account, the database can be accessed by any administrator on the system. In addition, any administrator who can log on to the system can run the server.

What to do next

To start the service, follow the instructions in Windows: Starting the server as a Windows service.

Windows

Windows: Starting the server as a Windows service

If you are running IBM Spectrum Protect™ on a Windows operating system, you can start the server as a service.

Before you begin

A Windows service must be created. The service was created automatically if you configured the server by using the configuration wizard. If the service was created automatically, you must configure the server to start as a service by following the steps in Windows: Configuring the server to start as a Windows service. Then, use this procedure to start the server as a service.

If you did not use the configuration wizard to create the service, you must create and configure the service manually. Follow the steps in Windows: Manually creating and configuring a Windows service.

Procedure

To start the server as a Windows service, complete the following steps:

1. Log on to the server with a user ID that is in the Administrators group.
2. From the Windows Start menu, click Run, type `services.msc`, and click OK.
3. In the Services window, select the server instance that you want to start, and click Start.

What to do next

Because the server service can issue requests that require action, it is important to monitor server activity with the Operations Center or the administrative client.

To view start and stop completion messages that are logged in the Windows application log, use the Event Viewer tool in the Administrative Tools folder.

Windows

Windows: Manually creating and configuring a Windows service

If you configured the server by using the configuration wizard, a Windows service was created automatically. If a service was not created automatically, you must create it.

Before you begin

To complete this procedure, you must log on with a user ID that is in the Administrators group.

Procedure

To create a Windows service and configure the startup options for the service, complete the following step:

Open a command window and enter the `sc.exe create` command:

```
sc.exe create server_name binPath= "path_to_server -k instance_name"
start= start_type obj= account_name password= password
```

where:

server_name

Specifies the name of the server service.

path_to_server

Specifies the path to the dsmsvc.exe executable file, including the file name. This path is the default path:

C:\Program Files\Tivoli\TSM\server

instance_name

Specifies the name of the DB2® instance, which is also the name of the server instance, for example, Server1.

start_type

Specifies the method for starting the service. To automatically start the service, enter `auto`. If you specify the `auto` option, the service starts automatically at system startup and restarts automatically whenever the system is restarted. To manually start the service, enter `demand`.

account_name

Specifies the user ID for the account under which the service runs. For example, the account name might be Administrator. This parameter is optional. If it is not specified, the Local System account is used.

password

Specifies the password for the *account_name* user account.

Tip: When you enter the command, ensure that you enter a space after each equal sign (=).

Results

The server starts as a Windows service.

Windows

Windows: Starting the server in the foreground

To directly interact with an IBM Spectrum Protect™ server, start the server in the foreground. For example, if you want to enter commands, start the server in the foreground.

Procedure

1. Change to the directory where the server is installed. For example, change to the `c:\program files\tivoli\tsm\server` directory.
2. Enter the following command:

```
dsmserv -k instance_name
```

where *instance_name* specifies the server instance.

Windows

Windows: Services associated with the server on Windows systems

When you start the IBM Spectrum Protect™ server as a service, other services start automatically. These services are associated with the database manager, DB2®.

The following services are associated with the server.

Service name	Purpose	Comments
TSM <i>Server_instance</i>	The service for the server instance that is named <i>Server_instance</i> . For example: TSM Server1	Set the start and stop options for this service to start and stop the server instance automatically. Each server instance runs as a separate service.

Service name	Purpose	Comments
DB2 - DB2TSM1 - <i>SERVER_INSTANCE</i>	The DB2 service for the server instance that is named <i>Server_instance</i> . For example: DB2 - DB2TSM1 - SERVER1	This service is automatically started when the service for the server instance is started. The DB2 service is not stopped automatically when you stop the service for the server. The system has one of these services for each server-instance service that is started on the system.
DB2 Governor (DB2TSM1)	A DB2 service that is created at installation time, and is required for all server instances.	Do not change the options for this service.
DB2 License Server (DB2TSM1)	A DB2 service that is created at installation time, and is required for all server instances.	Do not change the options for this service.
DB2 Management Server (DB2TSM1)	A DB2 service that is created at installation time, and is required for all server instances.	Do not change the options for this service.
DB2 Remote Command Server (DB2TSM1)	A DB2 service that is created at installation time, and is required for all server instances.	Do not change the options for this service.

Windows

Windows: Starting the server in maintenance mode

You can start the server in maintenance mode to avoid disruptions during maintenance and reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSEV utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dmserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dmserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode.

Operations that were disabled during maintenance mode are reenabled.

Windows: Stopping the server

You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.

About this task

To stop the server, issue the following command from the IBM Spectrum Protect™ command line:

```
halt
```

Windows: Registering licenses

Immediately register any IBM Spectrum Protect™ licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.

About this task

Use the REGISTER LICENSE command for this task. See REGISTER LICENSE for more details.

Example: Register a license

Register the base IBM Spectrum Protect license.

```
register license file=tsmbasic.lic
```

Windows: Preparing the server for database backup operations

To prepare the server for automatic and manual database backup operations, ensure that you specify a tape or file device class and complete other steps.

Procedure

1. Ensure that the IBM Spectrum Protect™ configuration is complete. If you did not use the configuration wizard (dsmicfgx) to configure the server, ensure that you completed the steps to manually configure the server for database backups.
2. Select the device class to be used for database backups, protect the master encryption key, and set a password. All of these actions are completed by issuing the SET DBRECOVERY command from the administrative command line:

```
set dbrecovery device_class_name protectkeys=yes password=password_name
```

where *device_class_name* specifies the device class to be used for database backup operations, and *password_name* specifies the password.

You must specify a device class name or the backup fails. By specifying PROTECTKEYS=YES, you ensure that the master encryption key is backed up during database backup operations.

Important: Create a strong password that is at least 8 characters long. Ensure that you remember this password. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database.

Example

To specify that database backups include a copy of the master encryption key for the server, run the following command:

```
set dbrecovery dbback protectkeys=yes password=protect8991
```

Windows: Running multiple server instances on a single system

You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

Multiply the memory and other system requirements for one server by the number of instances planned for the system.

Windows The set of files for one instance of the server is stored separately from the files used by another server instance on the same system. Use the steps in Windows: Creating the server instance for each new instance, optionally creating the new instance user.

To manage the system memory that is used by each server, use the DBMEMPERCENT server option to limit the percentage of system memory. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.

You can upgrade directly from V7.1 to V8.1. See the upgrade section (Upgrading to V8.1) for more details. When you upgrade and have multiple servers on your system, you must run the installation wizard only once. The installation wizard collects the database and variables information for all of your original server instances.

If you upgrade from IBM Spectrum Protect V6.3 to V8.1.5 and have multiple servers on your system, all instances that exist in DB2® V9.7 are dropped and recreated in DB2 V11.1. The wizard issues the `db2 upgrade db dbname` command for each database. The database environment variables for each instance on your system are also reconfigured during the upgrade process.

Windows A typical IBM Spectrum Protect installation involves one server instance on the IBM Spectrum Protect server computer. You might want to install a second instance if you are configuring in a clustered environment. You might also want to run more than one server on a large computer if you have multiple tape libraries or a disk-only configuration. After you install and configure the first IBM Spectrum Protect server, use the Server Initialization wizard to create additional IBM Spectrum Protect server instances on the same computer.

Windows By using the Server Initialization wizard, you can install up to four IBM Spectrum Protect server instances on a single system or cluster.

Related tasks:

[Running multiple server instances on a single system \(V7.1.1\)](#)

Windows: Monitoring the server

When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Procedure

1. Monitor the active log to ensure that the size is correct for the workload that is handled by the server instance.

When the server workload reaches its typical expected level, the space that is used by the active log is 80% - 90% of the space that is available to the active log directory. At that point, you might need to increase the amount of space. Whether you must increase the space depends on the types of transactions in the server workload. Transaction characteristics affect how the active log space is used.

The following transaction characteristics can affect the space usage in the active log:

- The number and size of files in backup operations
 - Clients such as file servers that back up large numbers of small files can cause large numbers of transactions that are completed quickly. The transactions might use a large amount of space in the active log, but for a short time.
 - Clients such as a mail server or a database server that back up large amounts of data in few transactions can cause small numbers of transactions that take a long time to complete. The transactions might use a small amount of space in the active log, but for a long time.
- Network connection types
 - Backup operations that occur over fast network connections cause transactions that complete more quickly. The transactions use space in the active log for a shorter time.
 - Backup operations that occur over relatively slower connections cause transactions that take a longer time to complete. The transactions use space in the active log for a longer time.

If the server is handling transactions with a wide variety of characteristics, the space that is used for the active log might increase and decrease significantly over time. For such a server, you might need to ensure that the active log typically has a smaller percentage of its space used. The extra space allows the active log to grow for transactions that take a long time to complete.

2. Monitor the archive log to ensure that space is always available.

Remember: If the archive log becomes full, and the failover archive log becomes full, the active log can become full, and the server stops. The goal is to make enough space available to the archive log so that it never uses all its available space.

You are likely to notice the following pattern:

- a. Initially, the archive log grows rapidly as typical client-backup operations occur.
- b. Database backups occur regularly, either as scheduled or done manually.
- c. After at least two full database backups occur, log pruning occurs automatically. The space that is used by the archive log decreases when the pruning occurs.
- d. Normal client operations continue, and the archive log grows again.
- e. Database backups occur regularly, and log pruning occurs as often as full database backups occur.

With this pattern, the archive log grows initially, decreases, and then might grow again. Over time, as normal operations continue, the amount of space that is used by the archive log should reach a relatively constant level.

If the archive log continues to grow, consider taking one or both of these actions:

- o Add space to the archive log. You might need to move the archive log to a different file system.
 - o Increase the frequency of full database backups, so that log pruning occurs more frequently.
3. If you defined a directory for the failover archive log, determine whether any logs get stored in that directory during normal operations. If the failover log space is being used, consider increasing the size of the archive log. The goal is that the failover archive log is used only under unusual conditions, not in normal operation.

Windows: Installing an IBM Spectrum Protect server fix pack

IBM Spectrum Protect™ maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.

Before you begin

To install a fix pack or interim fix to the server, install the server at the level on which you want to run it. You do not have to start the server installation at the base release level. For example, if you currently have V8.1.1 installed, you can go directly to the latest fix pack for V8.1. You do not have to start with the V8.1.0 installation if a maintenance update is available.

You must have the IBM Spectrum Protect license package installed. The license package is provided with the purchase of a base release. When you download a fix pack or interim fix from Fix Central, install the server license that is available on the Passport Advantage® website. To display messages and help in a language other than US English, install the language package of your choice.

If you upgrade the server to V8.1.5 or later, and then revert the server to a level that is earlier than V8.1.5, you must restore the database to a point in time before the upgrade. During the upgrade process, complete the required steps to ensure that the database can be restored: back up the database, the volume history file, the device configuration file, and the server options file. For more information, see Windows: Reverting from Version 8.1.5 to a previous server.

If you are using the client management service, ensure that you upgrade it to the same version as the IBM Spectrum Protect server.

Ensure that you retain the installation media from the base release of the installed server. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Visit the IBM® Support Portal for the following information:

- A list of the latest maintenance and download fixes. Click **Downloads** and apply any applicable fixes.
- Details about obtaining a base license package. Search for **Downloads > Passport Advantage**.
- Supported platforms and system requirements. Search for **IBM Spectrum Protect supported operating systems**.

Ensure that you upgrade the server before you upgrade backup-archive clients. If you do not upgrade the server first, communication between the server and clients might be interrupted.

Attention: Do not alter the DB2® software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Procedure

To install a fix pack or interim fix, complete the following steps:

1. Back up the database. The preferred method is to use a snapshot backup. A snapshot backup is a full database backup that does not interrupt any scheduled database backups. For example, issue the following IBM Spectrum Protect administrative command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information. Issue the following IBM Spectrum Protect administrative command:

```
backup devconfig filenames=file_name
```

where *file_name* specifies the name of the file in which to store device configuration information.

3. Save the volume history file to another directory or rename the file. Issue the following IBM Spectrum Protect administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named dsmserv.opt. The file is in the server instance directory.
5. Halt the server before installing a fix pack or interim fix. Use the HALT command.
6. Ensure that extra space is available in the installation directory. The installation of this fix pack might require additional temporary disk space in the installation directory of the server. The amount of additional disk space can be as much as that required for installing a new database as part of an IBM Spectrum Protect installation. The IBM Spectrum Protect installation wizard displays the amount of space that is required for installing the fix pack and the available amount. If the required amount of space is greater than the available amount, the installation stops. If the installation stops, add the required disk space to the file system and restart the installation.
7. Obtain the package file for the fix pack or interim fix that you want to install from the IBM Support Portal, Passport Advantage, or Fix Central.
8. **Windows** Change to the directory where you placed the executable file. Then, either double-click the following executable file or enter the following command on the command line to extract the installation files.
Tip: The files are extracted to the current directory. Ensure that the executable file is in the directory where you want the extracted files to be located.

```
8.x.x.x-IBM-SPSRV-platform.exe
```

where: *platform* denotes the operating system that IBM Spectrum Protect is to be installed on.

9. Select one of the following ways of installing IBM Spectrum Protect.

Important: After a fix pack is installed, it is not necessary to go through the configuration again. You can stop after completing the installation, fix any errors, then restart your servers.

Install the IBM Spectrum Protect software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

Windows: Installing IBM Spectrum Protect by using the installation wizard

Tip: After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.

Command line in console mode

Follow the instructions for your operating system:

Windows: Installing IBM Spectrum Protect by using console mode

Silent mode

Follow the instructions for your operating system:

Windows: Installing IBM Spectrum Protect in silent mode

Tip: If you have multiple server instances on your system, run the installation wizard only once. The installation wizard upgrades all server instances.

Results

Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click File > View Log. To collect log files, from the IBM Installation Manager tool, click Help > Export Data for Problem Analysis.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

- **Windows** C:\ProgramData\IBM\Installation Manager\logs
- **Windows** Windows: Applying a fix pack to IBM Spectrum Protect 8.1.5 in a clustered environment for Windows
To take advantage of new product features, you can upgrade a server that is installed on a Windows operating system in a clustered environment from V6.3 or V7.1 to IBM Spectrum Protect V8.1.5.

Windows: Reverting from Version 8.1.5 to a previous server

If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect™ server with minimal loss of data.

Before you begin

You must have the following items from the earlier version of the server:

- Server database backup
- Volume history file
- Device configuration file
- Server options file

About this task

Use the same instructions whether you are reverting within releases or to an earlier release, for example, from 8.1.3 to 8.1.2 or from 8.1.3 to 7.1.2. The older version must match the version that you used before the upgrade to V8.1.

Attention: Specify the REUSEDELAY parameter to help prevent backup-archive client data loss when you revert the server to a previous version.

- **Windows** Windows: Reverting to the previous server version in a cluster configuration
If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect server with minimal loss of data.

Steps for reverting to the previous server version

About this task

Complete the following steps on the system that has the V8.1 server.

Procedure

1. Halt the server to shut down all server operations by using the HALT command.
2. Remove the database from the database manager, then delete the database and recovery log directories.
 - a. Manually remove the database. One way to remove it is by issuing this command: **Windows**

```
dsmserv -k instance_name removedb tsmdb1
```
 - b. If you must reuse the space that is occupied by the database and recovery log directories, you can now delete these directories.
3. Use the uninstallation program to uninstall the V8.1 server. Uninstallation removes the server and the database manager, with their directories. For details, see Windows: Uninstalling IBM Spectrum Protect.

4. Stop the cluster service. Reinstall the version of the server program that you were using before the upgrade to V8.1.5. This version must match the version that your server was running when you created the database backup that you restore in a later step. For example, the server was at V7.1.7 before the upgrade, and you intend to use the database backup that was in use on this server. You must install the V7.1.7 fix pack to be able to restore the database backup.
5. Configure the new server database by using the configuration wizard. To start the wizard, issue the following command:

Windows

```
/dsmicfgx
```

6. Ensure that no servers are running in the background.
7. Restore the database to a point in time before the upgrade.
8. Copy the following files to the instance directory.
 - o Device configuration file
 - o Volume history file
 - o The server options file (typically dsmserv.opt)
9. If you enabled data deduplication for any FILE-type storage pools that existed before the upgrade, or if you moved data that existed before the upgrade into new storage pools while using the V8.1.5 server, you must complete additional recovery steps. For more details, see Additional recovery steps if you created new storage pools or enabled data deduplication.
10. If the REUSEDelay parameter setting on storage pools is less than the age of the database that you restored, restore volumes on any sequential-access storage pools that were reclaimed after that database backup. Use the RESTORE VOLUME command.
If you do not have a backup of a storage pool, audit the reclaimed volumes by using the AUDIT VOLUME command, with the FIX=YES parameter to resolve inconsistencies. For example:

```
audit volume volume_name fix=yes
```
11. If client backup or archive operations were completed using the V8.1 server, audit the storage pool volumes on which the data was stored.

Additional recovery steps if you created new storage pools or enabled data deduplication

If you created new storage pools, turned on data deduplication for any FILE-type storage pools, or did both while your server was running as a V8.1.5 server, you must complete more steps to return to the previous server version.

Before you begin

To complete this task, you must have a complete backup of the storage pool that was created before the upgrade to V8.1.5.

About this task

Use this information if you did either or both of the following actions while your server was running as a V8.1.5 server:

- You enabled the data deduplication function for any storage pools that existed before the upgrade to V8.1.5 program. Data deduplication applies only to storage pools that use a FILE device type.
- You created new primary storage pools after the upgrade *and* moved data that was stored in other storage pools into the new storage pools.

Complete these steps after the server is again restored to V7.

Procedure

- For each storage pool for which you enabled the data deduplication function, restore the entire storage pool by using the RESTORE STGPOOL command.
- For storage pools that you created after the upgrade, determine what action to take. Data that was moved from existing V8 storage pools into the new storage pools might be lost because the new storage pools no longer exist in your restored V8 server. Possible recovery depends on the type of storage pool:
 - o If data was moved from V8 DISK-type storage pools into a new storage pool, space that was occupied by the data that was moved was probably reused. Therefore, you must restore the original V8 storage pools by using the storage pool backups that were created before the upgrade to V8.1.5.

If *no* data was moved from V8 DISK-type storage pools into a new storage pool, then audit the storage pool volumes in these DISK-type storage pools.

- o If data was moved from V8 sequential-access storage pools into a new storage pool, that data might still exist and be usable in storage pool volumes on the restored V8 server. The data might be usable if the REUSEDELAY parameter for the storage pool was set to a value that prevented reclamation while the server was running as a V8.1.5 server. If any volumes were reclaimed while the server was running as a V8.1.5 server, restore those volumes from storage pool backups that were created before the upgrade to V8.1.5.

Windows: Reference: DB2 commands for IBM Spectrum Protect server databases

Use this list as reference when you are directed to issue DB2® commands by IBM® support.

Purpose




After using the wizards to install and configure IBM Spectrum Protect™, you seldom need to issue DB2 commands. A limited set of DB2 commands that you might use or be asked to issue are listed in Table 1. This list is supplemental material only and is not a comprehensive list. There is no implication that an IBM Spectrum Protect administrator will use it on a daily or ongoing basis. Samples of some commands are provided. Details of output are not listed.

For a full explanation of the commands described here and of their syntax, see the DB2 product information.

Table 1. DB2 commands

Command	Description	Example
<small>Windows</small> db2cmd	<small>Windows</small> Opens the command line processor DB2 window, and initializes the DB2 command-line environment.	<small>Windows</small> Open the DB2 command window: db2cmd
db2icrt	Creates DB2 instances in the home directory of the instance owner. Tip: The IBM Spectrum Protect configuration wizard creates the instance used by the server and database. After a server is installed and configured through the configuration wizard, the db2icrt command is generally not used. <small>Windows</small> This utility is located in the DB2PATH\bin directory where DB2PATH is the location where the DB2 copy is installed.	Manually create an IBM Spectrum Protect instance. Enter the command on one line: <pre>/opt/tivoli /tsm/db2/in stance/ db2icrt -a server -u instance_na me instance_na me</pre>
db2set	Displays DB2 variables.	List DB2 variables: db2set
CATALOG DATABASE	Stores database location information in the system database directory. The database can be located either on the local workstation or on a remote database partition server. The server configuration wizard takes care of any catalog needed for using the server database. Run this command manually, after a server is configured and running, only if something in the environment changes or is damaged.	Catalog the database: db2 catalog database tsmdb1

Command	Description	Example
CONNECT TO DATABASE	Connects to a specified database for command-line interface (CLI) use.	Connect to the IBM Spectrum Protect database from a DB2 CLI: db2 connect to tsmdb1
GET DATABASE CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Show the configuration information for a database alias: db2 get dbcfg for tsmdb1 Retrieve information in order to verify settings such as database configuration, log mode, and maintenance. db2 get dbconfig for tsmdb1 show detail
GET DATABASE MANAGER CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Retrieve configuration information for the database manager: db2 get dbmcfg
GET HEALTH SNAPSHOT	Retrieves the health status information for the database manager and its databases. The information returned represents a snapshot of the health state at the time the command was issued. IBM Spectrum Protect monitors the state of the database using the health snapshot and other mechanisms that are provided by DB2. There might be cases where the health snapshot or other DB2 documentation indicates that an item or database resource might be in an alert state. Such a case indicates that action must be considered to remedy the situation. IBM Spectrum Protect monitors the condition and responds appropriately. Not all declared alerts by the DB2 database are acted on.	Receive a report on DB2 health monitor indicators: db2 get health snapshot for database on tsmdb1

Command	Description	Example
GRANT (Database Authorities)	Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.	Grant access to the user ID itmuser: db2 GRANT CONNECT ON DATABASE TO USER itmuser db2 GRANT CREATETAB ON DATABASE TO USER itmuser
RUNSTATS	Updates statistics about the characteristics of a table and associated indexes or statistical views. These characteristics include number of records, number of pages, and average record length. To see a table, issue this utility after updating or reorganizing the table. A view must be enabled for optimization before its statistics can be used to optimize a query. A view that is enabled for optimization is known as a statistical view. Use the DB2 ALTER VIEW statement to enable a view for optimization. Issue the RUNSTATS utility when changes to underlying tables substantially affect the rows returned by the view. Tip: The server configures DB2 to run the RUNSTATS command as needed.	Update statistics on a single table. db2 runstats on table SCHEMA_NAME .TABLE_NAME with distribution and sampled detailed indexes all
 setdb2instance	 Determines which instance applies to the current session.	 Determine which instance is applicable: set db2instance =tsminst1
SET SCHEMA	Changes the value of the CURRENT SCHEMA special register, in preparation for issuing SQL commands directly through the DB2 CLI. Tip: A special register is a storage area that is defined for an application process by the database manager. It is used to store information that can be referenced in SQL statements.	Set the schema for IBM Spectrum Protect: db2 set schema tsmdb1
START DATABASE MANAGER	Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts. Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Start the database manager: db2start

Command	Description	Example
STOP DATA BASE MAN AGE R	<p>Stops the current database manager instance. Unless explicitly stopped, the database manager continues to be active. This command does not stop the database manager instance if any applications are connected to databases. If there are no database connections, but there are instance attachments, the command forces the instance attachments to stop first. Then, it stops the database manager. This command also deactivates any outstanding database activations before stopping the database manager.</p> <p>This command is not valid on a client.</p> <p>The server starts and stops the instance and database whenever the server starts and halts.</p> <p>Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.</p>	<p>Stop the database manager:</p> <pre>db2 stop dbm</pre>

Windows: Uninstalling IBM Spectrum Protect

You can use the following procedures to uninstall IBM Spectrum Protect™. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Before you begin

Complete the following steps before you uninstall IBM Spectrum Protect:

- Complete a full database backup.
- Save a copy of the volume history and device configuration files.
- Store the output volumes in a safe location.

Windows Attention: Do not use the Add/Remove Programs tool in the Windows Control Panel to uninstall IBM Spectrum Protect. Use only the uninstallation procedure that is described in this section.

About this task

You can uninstall IBM Spectrum Protect by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

- **Windows: Uninstalling IBM Spectrum Protect by using a graphical wizard**
You can uninstall IBM Spectrum Protect by using the IBM® Installation Manager installation wizard.
- **Windows: Uninstalling IBM Spectrum Protect in console mode**
To uninstall IBM Spectrum Protect by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.
- **Windows: Uninstalling IBM Spectrum Protect in silent mode**
To uninstall IBM Spectrum Protect in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.
- **Windows: Uninstalling and reinstalling IBM Spectrum Protect**
If you plan to manually reinstall IBM Spectrum Protect instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.
- **Windows: Uninstalling IBM Installation Manager**
You can uninstall IBM Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

What to do next

See Windows: Installing the server components for installation steps to reinstall the IBM Spectrum Protect components.

Windows: Uninstalling IBM Spectrum Protect by using a graphical wizard

You can uninstall IBM Spectrum Protect™ by using the IBM® Installation Manager installation wizard.

Procedure

1. Start the Installation Manager.

Windows Open the Installation Manager from the Start menu.

2. Click Uninstall.
3. Select IBM Spectrum Protect server, and click Next.
4. Click Uninstall.
5. Click Finish.

Windows: Uninstalling IBM Spectrum Protect in console mode

To uninstall IBM Spectrum Protect™ by using the command line, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **Windows** eclipse\tools

For example:

- o **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools
2. From the tools directory, issue the following command:
 - o **Windows** imcl.exe -c
 3. To uninstall, enter 5.
 4. Choose to uninstall from the IBM Spectrum Protect package group.
 5. Enter N for Next.
 6. Choose to uninstall the IBM Spectrum Protect server package.
 7. Enter N for Next.
 8. Enter U for Uninstall.
 9. Enter F for Finish.

Windows: Uninstalling IBM Spectrum Protect in silent mode

To uninstall IBM Spectrum Protect™ in silent mode, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the IBM Spectrum Protect server components. IBM Spectrum Protect includes a sample response file, `uninstall_response_sample.xml`, in the input directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

If you want to uninstall all IBM Spectrum Protect components, leave `modify="false"` set for each component in the response file. If you do not want to uninstall a component, set the value to `modify="true"`.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **Windows** eclipse\tools

For example:

- o **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools
2. From the tools directory, issue the following command, where `response_file` represents the response file path, including the file name:
 - o **Windows**

```
imcl.exe -input response_file -silent
```

The following command is an example:

Windows

```
imcl.exe -input C:\tmp\input\uninstall_response.xml -silent
```

Windows: Uninstalling and reinstalling IBM Spectrum Protect

If you plan to manually reinstall IBM Spectrum Protect™ instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.

About this task

To manually uninstall and reinstall IBM Spectrum Protect, complete the following steps:

1. **Windows** Make a list of your current server instances before proceeding to the uninstallation. Run the following command:

```
db2ilist
```

2. Run the following commands for every server instance:

Windows

```
db2 attach to server1
db2 get dbm cfg show detail
db2 detach
```

Keep a record of the database path for each instance.

3. Uninstall IBM Spectrum Protect. See [Windows: Uninstalling IBM Spectrum Protect](#).

Windows

After uninstalling IBM Spectrum Protect, check the Control Panel > Add or Remove Programs to verify that IBM Spectrum Protect DB2® is uninstalled.

4. When you uninstall any supported version of IBM Spectrum Protect, including a fix pack, an instance file is created. The instance file is created to help reinstall IBM Spectrum Protect. Check this file and use the information when you are prompted for the instance credentials when reinstalling. In silent installation mode, you provide these credentials using the `INSTANCE_CRED` variable.

You can find the instance file in the following location:

- o **Windows** C:\ProgramData\IBM\Tivoli\TSM\instanceList.obj in the IBM Spectrum Protect server installation directory

5. Reinstall IBM Spectrum Protect. See [Windows: Installing the server components](#).

If the `instanceList.obj` file does not exist, you need to recreate your server instances using the following steps:

- a. Recreate your server instances. See [Windows: Creating the server instance](#).
Tip: The installation wizard configures the server instances but you must verify that they exist. If they do not exist, you must manually configure them.
- b. Catalog the database. Log in to each server instance as the instance user, one at a time, and issue the following commands:

Windows

```
set db2instance=server1
db2 catalog database tsmdb1
db2 attach to server1
db2 update dbm cfg using dftdbpath instance_drive
db2 detach
```

- c. Verify that IBM Spectrum Protect recognizes the server instance by listing your directories. Your home directory appears if you did not change it. Your instance directory does appear if you used the configuration wizard. Issue this command:

```
db2 list database directory
```

If you see TSMDB1 listed, you can start the server.

Windows: Uninstalling IBM Installation Manager

You can uninstall IBM® Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

Before you begin

Before you uninstall IBM Installation Manager, you must ensure that all packages that were installed by IBM Installation Manager are uninstalled. Close IBM Installation Manager before you start the uninstall process.

Windows To view installed packages, click Start > All Programs > IBM Installation Manager > View Installed Packages.

Procedure

To uninstall IBM Installation Manager, complete the following steps:

Windows

1. From the Start menu, click Control Panel > Programs and Features.
2. Select IBM Installation Manager and click Uninstall.

Upgrading to V8.1

To take advantage of new product features and updates, upgrade the IBM Spectrum Protect™ server to Version 8.1.5.

About this task

To upgrade the server on the same operating system, see the upgrade instructions. For instructions about migrating the server to a different operating system, see IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions.

Table 1. Upgrade instructions

To upgrade from this version	To this version	See this information
V8.1	V8.1 fix pack or interim fix	AIX Installing an IBM Spectrum Protect server fix pack Linux Installing an IBM Spectrum Protect server fix pack Windows Installing an IBM Spectrum Protect server fix pack
V7.1	V8.1	Installing the server and verifying the upgrade
V7.1	V8.1 fix pack or interim fix	AIX Installing an IBM Spectrum Protect server fix pack Linux Installing an IBM Spectrum Protect server fix pack Windows Installing an IBM Spectrum Protect server fix pack
V5.5, V6.2, or V6.3	V8.1	IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions

An upgrade from V7 to V8.1 takes approximately 20 - 50 minutes. Your environment might produce different results from the results that were obtained in the labs.

For information about upgrades in a clustered environment, see Upgrading the server in a clustered environment.

To revert to an earlier version of the server after an upgrade or migration, you must have a full database backup and the installation software for the original server. You must also have the following key configuration files:

- Volume history file
- Device configuration file
- Server options file
- Upgrading to V8.1
You can upgrade the server directly from V7.1 to V8.1. You do not have to uninstall V7.1.
- Upgrading the server in a clustered environment
To upgrade a server to V8.1.5 in a clustered environment, you must complete preparation and installation tasks. The

procedures vary, depending on the operating system and release.

Related information:

[IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions](#)

Upgrading to V8.1

You can upgrade the server directly from V7.1 to V8.1. You do not have to uninstall V7.1.

Before you begin

Ensure that you retain the installation media from the server base release that you are upgrading. If you installed the server components from a DVD, ensure that the DVD is available. If you installed the server components from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Tip: DVDs are no longer available with V8.1 and later.

Procedure

To upgrade the server to V8.1, complete the following tasks:

- **Planning the upgrade**
Before you upgrade the server from V7.1 to V8.1, you must review the relevant planning information, such as system requirements and release notes. Then, select an appropriate day and time to upgrade the system so that you can minimize the impact on production operations.
- **Preparing the system**
To prepare the system for the upgrade from V7.1 to V8.1, you must gather information about each DB2® instance. Then, back up the server database, save key configuration files, cancel sessions, and stop the server.
- **Installing the server and verifying the upgrade**
To complete the process of upgrading the server to V8.1, you must install the V8.1 server. Then, verify that the upgrade was successful by starting the server instance.

Planning the upgrade

Before you upgrade the server from V7.1 to V8.1, you must review the relevant planning information, such as system requirements and release notes. Then, select an appropriate day and time to upgrade the system so that you can minimize the impact on production operations.

About this task

In lab tests, the process of upgrading the server from V7.1 to V8.1 took 14 - 45 minutes. The results that you achieve might differ, depending on your hardware and software environment, and the size of the server database.

Procedure

1. Review the hardware and software requirements:

AIX System requirements for AIX® systems

Linux System requirements for Linux systems

Windows System requirements for Windows systems

For the latest updates related to system requirements, see the IBM Spectrum Protect™ support website at technote 1243309.

2. For special instructions or specific information for your operating system, review the Release notes for Version 8.1 server components and IBM Spectrum Protect server Version 8.1 fix pack readme files.
3. Select an appropriate day and time to upgrade your system to minimize the impact on production operations. The amount of time that is required to update the system depends on the database size and many other factors. When you start the upgrade process, clients cannot connect to the server until the new software is installed and any required licenses are registered again.

4. If you are upgrading the server from V6 or V7 to V8.1, verify that you have the system ID and password for the DB2 instance of the IBM Spectrum Protect server. These credentials are required to upgrade the system.

Preparing the system

To prepare the system for the upgrade from V7.1 to V8.1, you must gather information about each DB2® instance. Then, back up the server database, save key configuration files, cancel sessions, and stop the server.

Procedure

1. Log on to the computer where the server is installed.

AIX | **Linux** Ensure that you are logged on with the instance user ID.

Windows Ensure that you are logged on with the administrative user ID that was used to install the V7.1 server.

2. Obtain a list of DB2 instances. Issue the following system command:

AIX | **Linux**

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

Windows

```
db2ilist
```

The output might be similar to the following example:

AIX | **Linux**

```
tsminst1
```

Windows

```
SERVER1
```

Ensure that each instance corresponds to a server that is running on the system.

3. **AIX** | **Linux** For each DB2 instance, note the default database path, actual database path, database name, database alias, and any DB2 variables that are configured for the instance. Keep the record for future reference. This information is required to restore the V7.1 database.
4. **Windows** Gather information about each DB2 instance. Note the default database path, actual database path, database name, database alias, and any DB2 variables that are configured for the instance. Keep the record for future reference. This information is required to restore the V7.1 database.
 - a. Open the DB2 command window by issuing the following system command:

```
db2cmd
```

- b. To change the instance, issue the following system command:

```
set DB2INSTANCE=instance
```

where *instance* specifies the DB2 instance.

- c. Obtain the default database path for the DB2 instance by issuing the following system command:

```
db2 get dbm cfg | findstr DFTDBPATH
```

The output might be similar to the following example:

```
Default database path (DFTDBPATH) = D:
```

- d. Obtain information about the DB2 instance databases by issuing the following system command:

```
db2 list database directory
```

The output might be similar to the following example:

```
System Database Directory
```

```
Number of entries in the directory = 2
```

```
Database 1 entry:
```



```

Database alias                = TSMAL001
Database name                 = TSMDB1
Node name                     = TSMNODE1
Database release level       = d.00
Comment                       = TSM SERVER DATABASE VIA TCPIP
Directory entry type         = Remote
Catalog database partition number = -1
Alternate server hostname     =
Alternate server port number  =

```

Database 2 entry:

```

Database alias                = TSMDB1
Database name                 = TSMDB1
Local database directory      = D:
Database release level       = d.00
Comment                       =
Directory entry type         = Indirect
Catalog database partition number = 0
Alternate server hostname     =
Alternate server port number  =

```

- e. Obtain the DB2 instance variables by issuing the following system command:

```
db2set -all
```

The output might be similar to the following example:

```

[e] DB2CODEPAGE=1208
[e] DB2PATH=D:\TSM\db2
[i] DB2_PMODEL_SETTINGS=MAX_BACKGROUND_SYSAPPS:500
[i] DB2_SKIPINSERTED=ON
[i] DB2_KEEPTABLELOCK=OFF
[i] DB2_EVALUNCOMMITTED=ON
[i] DB2_VENDOR_INI=D:\Server1\tsmdbmgr.env
[i] DB2_SKIPDELETED=ON
[i] DB2INSTPROF=C:\ProgramData\IBM\DB2\DB2TSM1
[i] DB2COMM=TCPIP
[i] DB2CODEPAGE=819
[i] DB2_PARALLEL_IO=*
[g] DB2_EXTSECURITY=YES
[g] DB2_COMMON_APP_DATA_PATH=C:\ProgramData

[g] DB2PATH=D:\TSM\db2
[g] DB2INSTDEF=SERVER1

```

5. Connect to the server by using an administrative user ID.
6. Back up the database by using the BACKUP DB command. The preferred method is to create a snapshot backup, which is a full database backup that does not interrupt scheduled database backups. For example, you can create a snapshot backup by issuing the following command:

```
backup db type=dbsnapshot devclass=tapeclass
```
7. Back up the device configuration information to another directory by issuing the following administrative command:

```
backup devconfig filenames=file_name
```

where *file_name* specifies the name of the file in which to store device configuration information.
Tip: If you decide to restore the V7.1 database, this file is required.
8. Back up the volume history file to another directory. Issue the following administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.
Tip: If you decide to restore the V7.1 database, this file is required.
9. Save a copy of the server options file, which is typically named dsmserv.opt. The file is in the server instance directory.
10. Prevent activity on the server by disabling new sessions. Issue the following administrative commands:

```

disable sessions client
disable sessions server

```

11. Verify whether any sessions exist, and notify the users that the server will be stopped. To check for existing sessions, issue the following administrative command:

```
query session
```

12. Cancel sessions by issuing the following administrative command:

```
cancel session all
```

This command cancels all sessions except for your current session.

13. Stop the server by issuing the following administrative command:

```
halt
```

14. Verify that the server is shut down and no processes are running.

AIX | **Linux** Issue the following command:

```
ps -ef | grep dsmserv
```

Windows Open the Windows Task Manager application and review the list of active processes.

15. In the server instance directory of your installation, locate the NODELOCK file and move it to another directory, where you are saving configuration files. The NODELOCK file contains the previous licensing information for your installation. This licensing information is replaced when the upgrade is complete.

Related reference:

BACKUP DB (Back up the database)

BACKUP DEVCONFIG (Create backup copies of device configuration information)

BACKUP VOLHISTORY (Save sequential volume history information)

DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)

QUERY SESSION (Query client sessions)

CANCEL SESSION (Cancel one or more client sessions)

HALT (Shut down the server)

Installing the server and verifying the upgrade

To complete the process of upgrading the server to V8.1, you must install the V8.1 server. Then, verify that the upgrade was successful by starting the server instance.

Before you begin

AIX | **Linux** You must be logged on to the system by using the root user ID.

Windows You must be logged on to the system with the administrative user ID that was used to install the previous server.

You can obtain the installation package from an IBM® download site.

AIX | **Linux** Set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly.

1. To query the maximum file size value, run the following command:

```
ulimit -Hf
```

2. If the system user limit for maximum file size is not set to unlimited, change the setting to unlimited by completing the instructions in the documentation for your operating system.

About this task

By using the IBM Spectrum Protect™ installation software, you can install the following components:

- Server
Tip: The database (DB2®), the Global Security Kit (GSKit), and IBM Java™ Runtime Environment (JRE) are automatically installed when you select the server component.
- Server languages
- License

- Devices
- IBM Spectrum Protect for SAN
- Operations Center

Procedure

1. Download the appropriate package file from one of the following websites:
 - Download the server package from Passport Advantage® or Fix Central.
 - For the most recent information, updates, and maintenance fixes, go to the IBM Support Portal.
2. Complete the following steps:



- a. Verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document for your product.
 - IBM Spectrum Protect technote 4042944
 - IBM Spectrum Protect Extended Edition technote 4042945
 - IBM Spectrum Protect for Data Retention technote 4042946
- b. Download the package file to the directory of your choice. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.

Also, ensure that you have executable permission for the package file.

- c. If necessary, run the following command to change the file permissions:

```
chmod a+x package_name.bin
```

where *package_name* is like the following example:

AIX

```
8.1.x.000-IBM-SPSRV-AIX.bin
```

Linux

```
8.1.x.000-IBM-SPSRV-Linuxs390x.bin
8.1.x.000-IBM-SPSRV-Linuxx86_64.bin
8.1.x.000-IBM-SPSRV-Linuxppc64le.bin
```

In the examples, *8.1.x.000* represents the product release level.

- d. Extract the installation files by running the following command:

```
./package_name.bin
```

The package is large. Therefore, the extraction takes some time.



- a. Verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document for your product.
 - IBM Spectrum Protect technote 4042944
 - IBM Spectrum Protect Extended Edition technote 4042945
 - IBM Spectrum Protect for Data Retention technote 4042946
- b. Change to the directory where you placed the executable file.

Tip: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.
- c. To extract the installation files, double-click the executable file:

```
package_name.exe
```

Where *package_name* is similar to the following example:

```
8.1.x.000-SPSRV-WindowsX64.exe
```

The package is large. Therefore, the extraction takes some time.

3. **AIX** To ensure that the IBM Spectrum Protect wizards work correctly, verify that the following command is enabled:
lsuser

4. Install the IBM Spectrum Protect software by using one of the following methods. Install the IBM Spectrum Protect license during the installation process.

Tip: If you have multiple server instances on your system, install the IBM Spectrum Protect software only one time to upgrade all server instances.

Installation wizard

AIX To install the server by using the graphical wizard of IBM Installation Manager, follow the instructions in Installing IBM Spectrum Protect by using the installation wizard.

Linux To install the server by using the graphical wizard of IBM Installation Manager, follow the instructions in Installing IBM Spectrum Protect by using the installation wizard.

Windows To install the server by using the graphical wizard of IBM Installation Manager, follow the instructions in Installing IBM Spectrum Protect by using the installation wizard.

Ensure that your system meets the prerequisites for using the installation wizard. Then, complete the installation procedure. In the IBM Installation Manager window, click the Update or Modify icon.

Installing the server by using the console mode

AIX To install the server by using the console mode, follow the instructions in Installing Tivoli® Storage Manager by using console mode.

Linux To install the server by using the console mode, follow the instructions in Installing Tivoli Storage Manager by using console mode.

Windows To install the server by using the console mode, follow the instructions in Installing Tivoli Storage Manager by using console mode.

Review the information about installing the server in console mode and then complete the installation procedure.

Silent mode

AIX To install the server by using silent mode, follow the instructions in Installing Tivoli Storage Manager in silent mode.

Linux To install the server by using silent mode, follow the instructions in Installing Tivoli Storage Manager in silent mode.

Windows To install the server by using silent mode, follow the instructions in Installing Tivoli Storage Manager in silent mode.

Review the information about installing the server in silent mode and then complete the installation procedure.

After you install the software, you do not have to reconfigure the system.

5. Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click File > View Log. To collect log files, from the IBM Installation Manager tool, click Help > Export Data for Problem Analysis.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

- o **AIX** **Linux** /var/ibm/InstallationManager/logs
- o **Windows** C:\ProgramData\IBM\Installation Manager\logs

6. Go to the IBM Support Portal to obtain fixes. Click Fixes, updates, and drivers and apply any applicable fixes.

7. Verify that the upgrade was successful:

- a. Start the server instance.

AIX For instructions, see Starting the server instance.

Linux For instructions, see Starting the server instance.

- b. Monitor the messages that the server issues as it starts. Watch for error and warning messages, and resolve any issues.

- c. Verify that you can connect to the server by using the administrative client. To start an administrative client session, run the following IBM Spectrum Protect administrative command:

```
dsmadm c
```

- d. To obtain information about the upgraded system, run QUERY commands. For example, to obtain consolidated information about the system, run the following IBM Spectrum Protect administrative command:

```
query system
```

To obtain information about the database, run the following IBM Spectrum Protect administrative command:

```
query db format=detailed
```

8. **Windows** Verify that the upgrade was successful:

- a. Start the server instance. To start the server from the default directory, C:\Program Files\Tivoli\TSM, run the following IBM Spectrum Protect administrative command:

```
dsmserv -k server_instance
```

server_instance is the name of your server instance. Server1 is the default name for the first instance of the IBM Spectrum Protect server.

If you plan to run the server as a service under the Local System account, the Local System account must be explicitly granted access to the server database. For instructions, see Starting the server by using Windows services.

- b. Monitor the messages that the server issues as it starts. Watch for error and warning messages, and resolve any issues.
- c. Verify that you can connect to the server by using the administrative client. To start an administrative client session, run the following IBM Spectrum Protect administrative command:

```
dsmadm c
```

- d. To obtain information about the upgraded system, run QUERY commands. For example, to obtain consolidated information about the system, run the following IBM Spectrum Protect administrative command:

```
query system
```

To obtain information about the database, run the following IBM Spectrum Protect administrative command:

```
query db format=detailed
```

9. **AIX** **Linux** Register the licenses for the IBM Spectrum Protect server components that are installed on your system by running the REGISTER LICENSE command:

```
register license file=installation_directory/server/bin/component_name.lic
```

where *installation_directory* specifies the directory in which you installed the component, and *component_name* specifies the abbreviation for the component.

For example, if you installed the server in the default directory, /opt/tivoli/tsm, run the following command to register the license:

```
register license file=/opt/tivoli/tsm/server/bin/tsmbasic.lic
```

For example, if you installed IBM Spectrum Protect Extended Edition in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/tsmee.lic
```

For example, if you installed IBM Spectrum Protect for Data Retention in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/dataret.lic
```

Restriction:

You cannot use the IBM Spectrum Protect server to register licenses for the following products:

- o IBM Spectrum Protect for Mail
- o IBM Spectrum Protect for Databases
- o IBM Spectrum Protect for ERP

- o IBM Spectrum Protect for Space Management

The REGISTER LICENSE command does not apply to these licenses. The licensing for these products is done by IBM Spectrum Protect clients.

10. **Windows** Register the licenses for the server components that are installed on your system by running the REGISTER LICENSE command:

```
register license file=installation_directory\server\component_name.lic
```

Where *installation_directory* specifies the directory in which you installed the component, and *component_name* specifies the abbreviation for the component.

For example, if you installed the server in the default directory, c:\Program Files\Tivoli\TSM, run the following command to register the license:

```
register license file=c:\Program Files\Tivoli\TSM\server\tsmbasic.lic
```

For example, if you installed IBM Spectrum Protect Extended Edition in the c:\Program Files\Tivoli\TSM directory, run the following command:

```
register license file=c:\Program Files\Tivoli\TSM\server\tsmee.lic
```

For example, if you installed IBM Spectrum Protect for Data Retention in the c:\Program Files\Tivoli\TSM directory, run the following command:

```
register license file=c:\Program Files\Tivoli\TSM\server\dataret.lic
```

Restriction:

You cannot use the IBM Spectrum Protect server to register licenses for the following products:

- o IBM Spectrum Protect for Mail
- o IBM Spectrum Protect for Databases
- o IBM Spectrum Protect for ERP
- o IBM Spectrum Protect for Space Management

The REGISTER LICENSE command does not apply to these licenses. The licensing for these products is done by IBM Spectrum Protect clients.

11. Optional: To install an extra language package, use the modify function of the IBM Installation Manager.
12. Optional: To upgrade to a newer version of a language package, use the update function of the IBM Installation Manager.

What to do next

You can authenticate passwords with the LDAP directory server, or authenticate passwords with the IBM Spectrum Protect server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

Windows If a device driver is available on Windows for the tape drives or medium changers that you plan to use, use the device driver. If a device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Spectrum Protect device driver by running the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is C:\Program Files\Tivoli\TSM\device\drivers.

Related reference:

QUERY SYSTEM (Query the system configuration and capacity)

QUERY DB (Display database information)

REGISTER LICENSE (Register a new license)

AIX **Linux** **Windows**

Upgrading the server in a clustered environment

To upgrade a server to V8.1.5 in a clustered environment, you must complete preparation and installation tasks. The procedures vary, depending on the operating system and release.

Procedure

Follow the procedure for your operating system, source release, and target release:

AIX

Table 1. Procedures for upgrading the server in a clustered environment on an AIX operating system

Source release	Target release	Procedure
V8.1	V8.1.5 fix pack	Applying a fix pack to V8 in a clustered environment for AIX
V6.3 or V7.1	V8.1.5	Upgrading IBM Spectrum Protect from V6.3 or V7.1 to V8.1.5 in a clustered environment for AIX with a shared database instance Upgrading from V6.3 to V8.1.5 in a clustered environment for AIX with separate database instances
V5.5, V6.1, V6.2	V7.1.1 or later	IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions

Linux

Table 2. Procedures for upgrading the server in a clustered environment on a Linux operating system

Source release	Target release	Procedure
V6.3 or later	V8.1.5	Upgrading a server that is configured with System Automation for Multiplatforms

Windows

Table 3. Procedures for upgrading the server in a clustered environment on a Windows operating system

Source release	Target release	Procedure
V8.1	V8.1.5 fix pack	Applying a fix pack to V8 in a clustered environment for Windows
V6.3 or V7.1	V8.1.5	Upgrading V6.3 or V7.1 to V8.1 in a clustered environment on Windows
V5.5, V6.1, V6.2	V7.1 or later	IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions

- Upgrading IBM Spectrum Protect from V6.3 or V7.1 to V8.1.5 in a clustered environment for AIX with a shared database instance
You can upgrade an IBM Spectrum Protect server from or V6.3 or V7.1 to V8.1.5 in a clustered environment on AIX with a shared database instance. In this way, you can take advantage of the new features in IBM Spectrum Protect V8.1.5.
- Upgrading from V6.3 to V8.1.5 in a clustered environment for AIX with separate database instances
You can upgrade a server from V6.3 to V8.1.5 in a clustered environment on AIX with separate database instances. In this way, you can take advantage of the new features in V8.1.5.
- Upgrading IBM Spectrum Protect to V8.1.5 in a clustered environment for Linux
To take advantage of new features in IBM Spectrum Protect, you can upgrade the IBM Spectrum Protect server that is installed on a Linux operating system in a clustered environment.
- Upgrading a V6.3 or V7.1 server to V8.1.5 in a clustered environment for Windows
To take advantage of new product features, you can upgrade a server that is installed on a Windows operating system in a clustered environment from V6.3 or V7.1 to IBM Spectrum Protect V8.1.5.

AIX

Upgrading IBM Spectrum Protect from V6.3 or V7.1 to V8.1.5 in a clustered environment for AIX with a shared database instance

You can upgrade an IBM Spectrum Protect™ server from or V6.3 or V7.1 to V8.1.5 in a clustered environment on AIX® with a shared database instance. In this way, you can take advantage of the new features in IBM Spectrum Protect V8.1.5.

Before you begin

Ensure that you retain the installation media from the V6.3 or V7.1 server base release that you are upgrading. If you installed IBM Spectrum Protect from a DVD, ensure that the DVD is available. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, you must reinstall the license from the installation media of the server base release.

About this task

Use the following procedure when the DB2® instance directory is shared between the nodes in the cluster. The DB2 instance directory is in the following location:

If the DB2 instance directory is not shared between nodes, follow the instructions in Upgrading from V6.3 to V8.1.5 in a clustered environment for AIX with separate database instances.

Procedure

1. Back up the database by using the BACKUP DB command. The preferred method is to use a snapshot backup, which creates a full database backup without interrupting any scheduled backups. For example, you can create a snapshot backup by running the following command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information to another directory, by running the following command:

```
backup devconfig filenames=file_name
```

Where *file_name* specifies the name of the file in which to store device configuration information.

3. Back up the volume history file to another directory, by running the following command:

```
backup volhistory filenames=file_name
```

Where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named dsmserv.opt, which is in the server instance directory.
5. Stop all instances of the server. Verify that no server processes are running. If you are using application-level monitoring of the IBM Spectrum Protect server, use your clustering tool to suspend monitoring of the dsmserv application resource.
6. Verify that the database manager is not running for any instance. Determine whether any db2sysc processes are running. The owner of running processes indicates which instances are active. For each server instance owner, run the following command to stop DB2:

```
db2stop
```

7. On the primary node, install the IBM Spectrum Protect V8.1.5 server by running the ./install.sh command. For instructions, see Installing the server components. After you start the wizard, in the IBM Installation Manager window, click the Update or Modify icon.
8. Start each V8.1.5 server in the foreground:
 - a. Verify that you are logged in with the instance owner ID.
 - b. Navigate to the instance directory and run the following command:

```
/opt/tivoli/tsm/server/bin/dsmserv
```

Wait until you see the server prompt, which indicates that the server is started.

9. Stop the server for each IBM Spectrum Protect instance that is being upgraded. Issue the following command:

```
halt
```

Tip: Because the DB2 instance directory is shared between the nodes in the cluster, you do not have to move the shared resources to the secondary node in the cluster.

10. On each secondary node in the cluster, complete the following steps:
 - a. Install the IBM Spectrum Protect V8.1.5 server by running the ./install.sh command. For instructions, see Installing the server components. 8.1.
 - i. If you are running the installation wizard, in the IBM Installation Manager window, click the Update or Modify icon.
 - ii. If you are running the installation wizard, in the Instance Credentials panel, clear the Update this instance check box for each instance.
 - iii. If you are installing the server in console mode, at the prompt `Do you want update this instance?`, enter `NO` for each instance.
 - iv. If you are installing the server in silent mode, specify `FALSE` for the value of the `user.instance_name_update` variable for each instance.
 - b. Ensure that each IBM Spectrum Protect V8.1.5 server starts. If you are using application-level monitoring, use the clustering tool to start the server.

For instructions about starting the server, see Starting the server instance.

11. Register the licenses for the server components that are installed on your system by running the REGISTER LICENSE command:


```
register license file=installation_directory/server/bin/component_name.lic
```

Where *installation_directory* specifies the directory in which you installed the component and *component_name* specifies the abbreviation for the component.

For example, if you installed the server in the default directory, /opt/tivoli/tsm, run the following command to register the license:

```
register license file=/opt/tivoli/tsm/server/bin/tsmbasic.lic
```

For example, if you installed IBM Spectrum Protect Extended Edition in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/tsmee.lic
```

For example, if you installed IBM Spectrum Protect for Data Retention in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/dataret.lic
```

Restriction:

You cannot use the IBM Spectrum Protect server to register licenses for the following products:

- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for ERP
- IBM Spectrum Protect for Space Management

The REGISTER LICENSE command does not apply to these licenses. The licensing for these products is done by IBM Spectrum Protect clients.

Related reference:

BACKUP DB (Back up the database)

BACKUP DEVCONFIG (Create backup copies of device configuration information)

BACKUP VOLHISTORY (Save sequential volume history information)

HALT (Shut down the server)

REGISTER LICENSE (Register a new license)

AIX

Upgrading from V6.3 to V8.1.5 in a clustered environment for AIX with separate database instances

You can upgrade a server from V6.3 to V8.1.5 in a clustered environment on AIX® with separate database instances. In this way, you can take advantage of the new features in V8.1.5.

Before you begin

Ensure that you retain the installation media from the V6.3 or V7.1 server base release that you are upgrading. If you installed IBM Spectrum Protect™ from a DVD, ensure that the DVD is available. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, you must reinstall the license from the installation media of the server base release.

About this task

Use the following procedure when the DB2® instance directory is not shared between the nodes in the cluster. The DB2 instance directory is at the following location:

```
/home/tsminst1/sqlllib
```

If the DB2 instance directory is shared between the nodes in the cluster, follow the instructions in Upgrading IBM Spectrum Protect from V6.3 or V7.1 to V8.1.5 in a clustered environment for AIX with a shared database instance.

Procedure

1. Back up the database by using the BACKUP DB command. The preferred method is to use a snapshot backup, which creates a full database backup without interrupting any scheduled backups. For example, you can create a snapshot

backup by running the following command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information to another directory, by running the following command:

```
backup devconfig filenames=file_name
```

Where *file_name* specifies the name of the file in which to store device configuration information.

3. Back up the volume history file to another directory, by running the following command:

```
backup volhistory filenames=file_name
```

Where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named `dsmserv.opt`, which is in the server instance directory.
5. Stop all instances of the server. Verify that no server processes are running. If you are using application-level monitoring of the IBM Spectrum Protect server, use your clustering tool to suspend monitoring of the `dsmserv` application resource.
6. Verify that the database manager is not running for any instance. Determine whether any `db2sysc` processes are running. The owner of running processes indicates which instances are active. For each server instance owner, run the following command to stop DB2:

```
db2stop
```

7. Ensure that the shared resources for all IBM Spectrum Protect instances are on the primary node. Verify that no other nodes have write access to these resources during the upgrade. If the environment includes multiple instances of the server, shared resources for all instances must be accessible to the primary node.
8. On the primary node, install the V8.1.5 server by running the `./install.sh` command. For instructions, see *Installing the server components*. After you start the wizard, in the IBM Installation Manager window, click the Install icon; do not click the Update or Modify icon. To complete the upgrade from V6.3 to V8.1.5, you must install the V8.1.5 server.
9. Start each V8.1.5 server in the foreground:
 - a. Verify that you are logged in with the instance owner ID.
 - b. Navigate to the instance directory and run the following command:

```
/opt/tivoli/tsm/server/bin/dsmserv
```

Wait until you see the server prompt, which indicates that the server is started.

10. Stop the server for each IBM Spectrum Protect instance that is being upgraded. Run the following command:

```
halt
```

11. On each secondary node in the cluster, complete the following steps:
 - a. Move all shared resources to the secondary node. If the environment includes multiple instances of the server, shared resources for all instances must be accessible to the secondary nodes during the upgrade.
 - b. Stop all instances of the server. Verify that no server processes are running.
 - c. Verify that the database manager is not running for any instance. Determine whether any `db2sysc` processes are running. The owner of running processes indicates which instances are active. For each server instance owner, run the following command to stop DB2:

```
db2stop
```

- d. Install the V8.1.5 server by running the `./install.sh` command. For instructions, see *Installing the server components*.
 - i. If you are using the installation wizard, in the IBM Installation Manager window, click the Install icon; do not click the Update or Modify icon.
 - ii. If you are using the installation wizard, on the Instance Credentials page, select the Configure this instance on a secondary node of the cluster check box for each instance that you are configuring.
 - iii. If you are installing the server in console mode, at the prompt `Configure this instance on a secondary node of the cluster?`, enter `YES` for each instance.
 - iv. If you are installing the server in silent mode, specify `TRUE` for the value of the `user.instance_name_secondaryNode` variable for each instance.
- e. Ensure that each V8.1.5 server starts. If you are using application-level monitoring, use the clustering tool to start the server.

For instructions about starting the server, see *Starting the server instance*.

12. Register the licenses for the server components that are installed on your system by running the `REGISTER LICENSE` command:

```
register license file=installation_directory/server/bin/component_name.lic
```

Where *installation_directory* specifies the directory in which you installed the component and *component_name* specifies the abbreviation for the component.

For example, if you installed the server in the default directory, /opt/tivoli/tsm, run the following command to register the license:

```
register license file=/opt/tivoli/tsm/server/bin/tsmbasic.lic
```

For example, if you installed IBM Spectrum Protect Extended Edition in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/tsmee.lic
```

For example, if you installed IBM Spectrum Protect for Data Retention in the /opt/tivoli/tsm directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/dataret.lic
```

Restriction:

You cannot use the IBM Spectrum Protect server to register licenses for the following products:

- o IBM Spectrum Protect for Mail
- o IBM Spectrum Protect for Databases
- o IBM Spectrum Protect for ERP
- o IBM Spectrum Protect for Space Management

The REGISTER LICENSE command does not apply to these licenses. The licensing for these products is done by IBM Spectrum Protect clients.

Related reference:

BACKUP DB (Back up the database)

BACKUP DEVCONFIG (Create backup copies of device configuration information)

BACKUP VOLHISTORY (Save sequential volume history information)

HALT (Shut down the server)

REGISTER LICENSE (Register a new license)

Linux

Upgrading IBM Spectrum Protect to V8.1.5 in a clustered environment for Linux

To take advantage of new features in IBM Spectrum Protect™, you can upgrade the IBM Spectrum Protect server that is installed on a Linux operating system in a clustered environment.

Procedure

Follow the instructions in Configuring a Linux environment for clustering.

Windows

Upgrading a V6.3 or V7.1 server to V8.1.5 in a clustered environment for Windows

To take advantage of new product features, you can upgrade a server that is installed on a Windows operating system in a clustered environment from V6.3 or V7.1 to IBM Spectrum Protect™ V8.1.5.

Before you begin

Ensure that you retain the installation media from the V6.3 or V7.1 server base release that you are upgrading. If you installed the server from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, you must reinstall the license from the installation media of the server base release.

Procedure

1. Back up the database by using the BACKUP DB command. The preferred method is to use a snapshot backup, which provides a full database backup without interrupting scheduled backups. For example, you can run the following command to create a snapshot backup:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information to another directory. Run the following command:

```
backup devconfig filenames=file_name
```

Where *file_name* specifies the name of the file in which to store device configuration information.

3. Back up the volume history file to another directory. Run the following command:

```
backup volhistory filenames=file_name
```

Where *file_name* specifies the name of the file in which to store the volume history information.

4. Save a copy of the server options file, typically named dmserv.opt, which is in the server instance directory.
5. Ensure that the resource group is on the primary node, and that all nodes in the cluster are running. Take the following actions on the primary node:
 - a. In the Failover Cluster Manager window, take the server resource offline and delete it:
 - i. Select Services and applications, and then select the cluster group. The server resource is displayed in the Other Resources section.
 - ii. Select the server resource, and click Take this resource offline.
 - iii. To delete the server resource, select it, and click Delete.
 - b. In the Failover Cluster Manager window, delete the network name and IP address:
 - i. In the Server name section, expand the network name to view the IP address. Note the network name and IP address.
 - ii. Select the network name and the IP address, and click Remove.
 - c. In the Failover Cluster Manager window, take the DB2® server resource offline:
 - i. Select Services and applications, and then select the cluster group. The IBM Spectrum Protect server resource is displayed in the Other Resources section.
 - ii. Select a DB2 server resource, for example, SERVER1, and click Take this resource offline.
6. Ensure that the server is running on the primary node. Complete the following steps on all other cluster nodes:
 - a. Install the IBM Spectrum Protect V 8.1.5 server.
 - b. Stop the cluster service. One way to stop it is by using the Services Application. Right-click Cluster Service and select Stop.
 - c. Delete the tsmsvrrscexX64.dll and tsmsvrrscx64.dll files from the C:\Windows\Cluster directory.
 - d. Copy the following DLL files from the installation directory to the C:\Windows\Cluster directory:
 - tsmsvrrscexX64.dll
 - tsmsvrrscx64.dll
 - e. Copy the following DLL file from the installation directory to the C:\TSM\db2\security\plugin\IBM\server directory: dsmdb2pw64.dll
 - f. Start the cluster service. One way to start it is by using the Services Application. Right-click Cluster Service and select Start.
7. In the Failover Cluster Manager, move the IBM Spectrum Protect server instance from the primary node to another node in the cluster.
8. Complete the following steps on the primary node:
 - a. Install the IBM Spectrum Protect V 8.1.5 server.
 - b. Stop the cluster service.
 - c. Delete the tsmsvrrscexX64.dll and tsmsvrrscx64.dll files from the C:\Windows\Cluster directory.
 - d. Copy the following DLL files from the installation directory to the C:\Windows\Cluster directory:
 - tsmsvrrscexX64.dll
 - tsmsvrrscx64.dll
 - e. Copy the following DLL file from the installation directory to the C:\TSM\db2\security\plugin\IBM\server directory: dsmdb2pw64.dll
 - f. Start the cluster service.
9. Optional: Move the IBM Spectrum Protect server instance back to the primary node.

What to do next

If a device driver is available on Windows for the tape drives or medium changers that you plan to use, use the device driver. If a device driver is not available, install the IBM Spectrum Protect device driver by running the dpinst.exe /a command. The dpinst.exe file is in the device driver directory, and the default location is C:\Program Files\Tivoli\TSM\device\drivers.

Related reference:

BACKUP DB (Back up the database)
 BACKUP DEVCONFIG (Create backup copies of device configuration information)
 BACKUP VOLHISTORY (Save sequential volume history information)
 REGISTER LICENSE (Register a new license)

AIX Linux Windows

Installing and upgrading the Operations Center

The IBM Spectrum Protect™ Operations Center is the web-based interface for managing your storage environment.

Before you begin

Before you install and configure the Operations Center, review the following information:

- System requirements for the Operations Center
 - Operations Center computer requirements
 - Hub and spoke server requirements
 - Operating system requirements
 - Web browser requirements
 - Language requirements
 - Requirements and limitations for IBM Spectrum Protect client management services
- Administrator IDs that the Operations Center requires
- IBM Installation Manager
- Installation checklist
- Obtaining the Operations Center installation package

About this task

Table 1 lists the methods for installing or uninstalling the Operations Center and indicates where to find the associated instructions.

For information about upgrading the Operations Center, see [Upgrading the Operations Center](#).

Table 1. Methods for installing or uninstalling the Operations Center

Method	Instructions
Graphical wizard	<ul style="list-style-type: none"> • Installing the Operations Center by using a graphical wizard • Uninstalling the Operations Center by using a graphical wizard
Console mode	<ul style="list-style-type: none"> • Installing the Operations Center in console mode • Uninstalling the Operations Center in console mode
Silent mode	<ul style="list-style-type: none"> • Installing the Operations Center in silent mode • Uninstalling the Operations Center in silent mode

- Planning to install the Operations Center
 Before you install the Operations Center, you must understand the system requirements, the administrator IDs that the Operations Center requires, and the information that you must provide to the installation program.
- Installing the Operations Center
 You can install the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.
- Upgrading the Operations Center
 You can upgrade the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.
- Getting started with the Operations Center
 Before you can use the Operations Center to manage your storage environment, you must configure it.
- [AIX Linux Troubleshooting the Operations Center installation](#)
 If a problem occurs with the Operations Center installation and you cannot solve it, you can consult the descriptions of known problems for a possible solution.

- Uninstalling the Operations Center
You can uninstall the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.
- Rolling back to a previous version of the Operations Center
By default, IBM Installation Manager saves earlier versions of a package to roll back to if you experience a problem with later versions of updates, fixes, or packages.

AIX

Linux

Windows

Planning to install the Operations Center

Before you install the Operations Center, you must understand the system requirements, the administrator IDs that the Operations Center requires, and the information that you must provide to the installation program.

About this task

From the Operations Center, you can manage the following primary aspects of the storage environment:

- IBM Spectrum Protect™ servers and clients
- Services such as backup and restore, archive and retrieve, and migrate and recall
- Storage pools and storage devices

The Operations Center includes the following features:

User interface for multiple servers

You can use the Operations Center to manage one or more IBM Spectrum Protect servers.

In an environment with multiple servers, you can designate one server as a *hub server* and the others as *spoke servers*. The hub server can receive alerts and status information from the spoke servers and present the information in a consolidated view in the Operations Center.

Alert monitoring

An *alert* is a notification of a relevant problem on the server and is triggered by a server message. You can define which server messages trigger alerts, and only those messages are reported as alerts in the Operations Center or in an email.

This alert monitoring can help you identify and track relevant problems on the server.

Convenient command-line interface

The Operations Center includes a command-line interface for advanced features and configuration.

- System requirements for the Operations Center
Before you install the Operations Center, ensure that your system meets the minimum requirements.
- Administrator IDs that the Operations Center requires
An administrator must have a valid ID and password on the hub server to log in to the Operations Center. An administrator ID is also assigned to the Operations Center so that the Operations Center can monitor servers.
- IBM Installation Manager
The Operations Center uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.
- Installation checklist
Before you install the Operations Center, you must verify certain information, such as the installation credentials, and you must determine the input to provide to IBM Installation Manager for the installation.

AIX

Linux

Windows

System requirements for the Operations Center

Before you install the Operations Center, ensure that your system meets the minimum requirements.

Use the Operations Center System Requirements Calculator to estimate the system requirements for running the Operations Center and the hub and spoke servers that are monitored by the Operations Center.

Requirements that are verified during the installation

Table 1 lists the prerequisite requirements that are verified during the installation and indicates where to find more information about these requirements.

Table 1. Requirements that are verified during the installation

Requirement	Details
Minimum memory requirement	Operations Center computer requirements
Operating system requirement	Operating system requirements
Host name for the computer where the Operations Center will be installed	Installation checklist
Requirements for the Operations Center installation directory	Installation checklist

- Operations Center computer requirements
You can install the Operations Center on a computer that is also running IBM Spectrum Protect server or on a different computer. If you install the Operations Center on the same computer as a server, that computer must meet the system requirements for both the Operations Center and the server.
- Hub and spoke server requirements
When you open the Operations Center for the first time, you must associate the Operations Center with one IBM Spectrum Protect server that is designated as the *hub server*. In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.
- Operating system requirements
The Operations Center is available for AIX, Linux, and Windows systems.
- Web browser requirements
The Operations Center can run in Apple, Google, Microsoft, and Mozilla web browsers.
- Language requirements
By default, the Operations Center uses the language that the web browser uses. However, the installation process uses the language that the operating system uses. Verify that the web browser and the operating system are set to the language that you require.
- Requirements and limitations for IBM Spectrum Protect client management services
IBM Spectrum Protect client management services is a component that you install on backup-archive clients to collect diagnostic information such as client log files. Before you install the client management service on your system, you must understand the requirements and limitations.

AIX | Linux | Windows

Operations Center computer requirements

You can install the Operations Center on a computer that is also running IBM Spectrum Protect™ server or on a different computer. If you install the Operations Center on the same computer as a server, that computer must meet the system requirements for both the Operations Center and the server.

Resource requirements

The following resources are required to run the Operations Center:

- One processor core
- 4 GB of memory
- 1 GB of disk space

The hub and spoke servers that are monitored by the Operations Center require additional resources, as described in Hub and spoke server requirements.

AIX | Linux | Windows

Hub and spoke server requirements

When you open the Operations Center for the first time, you must associate the Operations Center with one IBM Spectrum Protect™ server that is designated as the *hub server*. In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

If only one server is monitored by the Operations Center, that server is still called a hub server, even though no spoke servers are connected to it.

Table 1 indicates the version of IBM Spectrum Protect server that must be installed on the hub server and on each spoke server that is managed by the Operations Center.

Table 1. IBM Spectrum Protect server version requirements for hub and spoke servers

Operations Center	Version on the hub server	Version on each spoke server
V8.1.5	V8.1.5	V6.3.4 or later Restriction: Some Operations Center functions are not available for servers that use a version earlier than V8.1.5.

Number of spoke servers that a hub server can support

The number of spoke servers that a hub server can support depends on the configuration and on the version of IBM Spectrum Protect on each spoke server. However, a general guideline is that a hub server can support 10 - 20 V6.3.4 spoke servers but can support more V7.1 or later spoke servers.

- Tips for designing the hub and spoke server configuration
In designing the hub and spoke configuration, especially consider the resource requirements for status monitoring. Also, consider how you want to group hub and spoke servers and whether you want to use multiple hub servers.
- Tips for choosing a hub server
For the hub server, you must choose a server that has adequate resources and is located for minimal roundtrip network latency.

AIX

Linux

Windows

Tips for designing the hub and spoke server configuration

In designing the hub and spoke configuration, especially consider the resource requirements for status monitoring. Also, consider how you want to group hub and spoke servers and whether you want to use multiple hub servers.

Use the Operations Center System Requirements Calculator to estimate the system requirements for running the Operations Center and the hub and spoke servers that are monitored by the Operations Center.

Primary factors that affect performance

The following factors have the most significant impact on the performance of the Operations Center:

- The processor and memory on the computer where the Operations Center is installed
- The system resources of the hub and spoke servers, including the disk system that is in use for the hub server database
- The number of client nodes and virtual machine file spaces that are managed by the hub and spoke servers
- The frequency at which data is refreshed in the Operations Center

How to group hub and spoke servers

Consider grouping hub and spoke servers by geographic location. For example, managing the servers within the same data center can help prevent issues that are caused by firewalls or by inadequate network bandwidth between different locations. If necessary, you can further divide servers according to one or more of the following characteristics:

- The administrator who manages the servers
- The organizational entity that funds the servers
- Server operating system
- The language in which the servers run
Tip: If the hub and spoke servers are not running in the same language, you might see corrupted text in the Operations Center.

How to group hub and spoke servers in an enterprise configuration

In an enterprise configuration, a network of IBM Spectrum Protect™ servers are managed as a group. Changes that are made on the *configuration manager* can be distributed automatically to one or more *managed servers* in the network.

The Operations Center normally registers and maintains a dedicated administrator ID on the hub and spoke servers. This *monitoring administrator* must always have the same password on all the servers.

If you use an enterprise configuration, you can improve the process by which the administrator credentials are synchronized on spoke servers. To improve the performance and efficiency of maintaining the monitoring administrator ID, complete the following steps:

1. Designate the configuration manager server as the Operations Center hub server. During the hub server configuration, a monitoring administrator ID named `IBM-OC-hub_server_name` is registered.
2. On the hub server, add the monitoring administrator ID to a new or existing enterprise configuration profile. Issue the `NOTIFY SUBSCRIBERS` command to distribute the profile to the managed servers.
3. Add one or more of the managed servers as Operations Center spoke servers.

The Operations Center detects this configuration and allows the configuration manager to distribute and update the monitoring administrator ID on the spoke servers.

When to use multiple hub servers

If you have more than 10 - 20 V6.3.4 spoke servers, or if resource limitations require the environment to be partitioned, you can configure multiple hub servers, and connect a subset of the spoke servers to each hub server.

Restrictions:

- A single server cannot be both a hub server and a spoke server.
- Each spoke server can be assigned to only one hub server.
- Each hub server requires a separate instance of the Operations Center, each of which has a separate web address.

AIX

Linux

Windows

Tips for choosing a hub server

For the hub server, you must choose a server that has adequate resources and is located for minimal roundtrip network latency.

Attention: Do not use the same server as the hub server for multiple Operations Centers.

Use the following guidelines in deciding which server to designate as the hub server:

Choose a lightly loaded server

Consider a server that has a light load for operations such as client backup and archive. A lightly loaded server is also a good choice as the host system for the Operations Center.

Ensure that the server has the resources to handle both its typical server workload and the estimated workload for acting as the hub server.

Locate the server for minimal roundtrip network latency

Locate the hub server so that the network connection between the hub server and the spoke servers has a roundtrip latency that is no greater than 5 ms. This latency can typically be achieved when the servers are on the same local area network (LAN).

Networks that are poorly tuned, are heavily used by other applications, or have roundtrip latency much higher than 5 ms can degrade communications between the hub and spoke servers. For example, roundtrip latencies of 50 ms or higher can result in communication timeouts that cause spoke servers to disconnect or reconnect to the Operations Center. Such high latencies might be experienced in long-distance, wide area network (WAN) communications.

If spoke servers are a long distance from the hub server and experience frequent disconnects in the Operations Center, you can increase the value of the `ADMINCOMMTIMEOUT` option on each server to reduce the problem.

Verify that the hub server meets the resource requirements for status monitoring

Status monitoring requires extra resources on each server on which it is enabled. The resources that are required depend primarily on the number of clients that are managed by the hub and spoke servers. Fewer resources are used on a hub server with a V7.1 or later spoke server than on a hub server with a V6.3.4 spoke server.

Verify that the hub server meets the resource requirements for processor usage, database space, archive log space, and I/O operations per second (IOPS) capacity.

A hub server with high IOPS capacity can handle a larger amount of incoming status data from spoke servers. Use of the following storage devices for the hub server database can help meet this capacity:

- An enterprise-level solid-state drive (SSD)
- An external SAN disk storage device with multiple volumes or multiple spindles under each volume

In an environment with fewer than 1000 clients, consider establishing a baseline capacity of 1000 IOPS for the hub server database if the hub server manages any spoke servers.

Determine whether your environment requires multiple hub servers

If more than 10,000 - 20,000 client nodes and virtual machine file spaces are managed by one set of hub and spoke servers, the resource requirements might exceed what the hub server has available, especially if the spoke servers are V6.3.4 servers. Consider designating a second server as a hub server and moving spoke servers to the new hub server to balance the load.

AIX | Linux | Windows

Operating system requirements

The Operations Center is available for AIX®, Linux, and Windows systems.

You can run the Operations Center on the following systems:

- **AIX** AIX systems:
 - IBM® AIX V7.1 (64 bit) TL 4 and SP 2
 - IBM AIX V7.2 (64 bit) TL 0 and SP 2
- **Linux** Linux on x86_64 systems:
 - Red Hat Enterprise Linux 6.7
 - Red Hat Enterprise Linux 7.1
 - SUSE Linux Enterprise Server 11, Service Pack 4 or later
 - SUSE Linux Enterprise Server 12
- **Linux** Linux on System z (s390x 64-bit architecture) systems:
 - Red Hat Enterprise Linux 7.1
 - SUSE Linux Enterprise Server 12
- **Linux** Linux on Power Systems (little endian) systems:
 - Red Hat Enterprise Linux 7.3 with the PPC64LE architecture
- **Windows** Windows systems:
 - Microsoft Windows Server 2012: Standard, Enterprise, or Datacenter Edition (64-bit)
 - Microsoft Windows Server 2012 R2 (64-bit)
 - Microsoft Windows Server 2016

For the most up-to-date requirements information, see Software and Hardware Requirements.

AIX | Linux | Windows

Web browser requirements

The Operations Center can run in Apple, Google, Microsoft, and Mozilla web browsers.

For optimal viewing of the Operations Center in the web browser, ensure that the screen resolution for the system is set to a minimum of 1024 X 768 pixels.

For optimal performance, use a web browser that has good JavaScript performance, and enable browser caching.

The Operations Center can run in the following web browsers:

- Apple Safari on the iPad
Restriction: If Apple Safari is running on iOS 8.x or iOS 9.x, you cannot use a self-signed certificate for secure communication with the Operations Center without extra configuration of the certificate. Use a certificate authority (CA) certificate, or configure the self-signed certificate as needed. For instructions, see Technote <http://www.ibm.com/support/docview.wss?uid=swg21963153>.
- Google Chrome 54 or later
- Microsoft Internet Explorer 11 or later
- Mozilla Firefox ESR 45 or version 48 or later

Communication between the Operations Center and the web browser must be secured by using the Transport Layer Security (TLS) 1.2 protocol. The web browser must support TLS 1.2, and TLS 1.2 must be enabled. The web browser displays an SSL error if it does not meet these requirements.

AIX

Linux

Windows

Language requirements

By default, the Operations Center uses the language that the web browser uses. However, the installation process uses the language that the operating system uses. Verify that the web browser and the operating system are set to the language that you require.

AIX

Table 1. Operations Center language values that you can use on AIX® systems

Language	Language option value
Chinese, Simplified	zh_CN
Chinese, Simplified (UTF-8)	ZH_CN
Chinese, Traditional (Big5)	Zh_TW
Chinese, Traditional (UTF-8)	ZH_TW
Chinese, Traditional (euc_tw)	zh_TW
English	en_US
English (UTF-8)	EN_US
French	fr_FR
French (UTF-8)	FR_FR
German	de_DE
German (UTF-8)	DE_DE
Italian	it_IT
Italian (UTF-8)	IT_IT
Japanese (EUC)	ja_JP
Japanese (PC)	Ja_JP
Japanese (UTF-8)	JA_JP
Korean	ko_KR
Korean (UTF-8)	KO_KR
Portuguese, Brazilian	pt_BR
Portuguese, Brazilian (UTF-8)	PT_BR
Russian	ru_RU
Russian (UTF-8)	RU_RU
Spanish	es_ES
Spanish (UTF-8)	ES_ES

Linux

Table 2. Operations Center language values that you can use on Linux systems

Language	Language option value
Chinese, Simplified	zh_CN
Chinese, Simplified (GBK)	zh_CN.gb18030
Chinese, Simplified (UTF-8)	zh_CN.utf8

Language	Language option value
Chinese, Traditional (Big5)	Zh_TW
Chinese, Traditional (euc_tw)	zh_TW
Chinese, Traditional (UTF-8)	zh_TW.utf8
English, United States	en_US
English (UTF-8)	en_US.utf8
French	fr_FR
French (UTF-8)	fr_FR.utf8
German	de_DE
German (UTF-8)	de_DE.utf8
Italian	it_IT
Italian (UTF-8)	it_IT.utf8
Japanese (EUC)	ja_JP
Japanese (UTF-8)	ja_JP.utf8
Korean	ko_KR
Korean (UTF-8)	ko_KR.utf8
Portuguese, Brazilian	pt_BR
Portuguese, Brazilian (UTF-8)	pt_BR.utf8
Russian	ru_RU
Russian (UTF-8)	ru_RU.utf8
Spanish	es_ES
Spanish (UTF-8)	es_ES.utf8

Windows

Table 3. Operations Center language values that you can use on Windows systems

Language	Language option value
Chinese, Simplified	chs
Chinese, Traditional	cht
English	ameng
French	fra
German	deu
Italian	ita
Japanese (Shift-JIS)	jpn
Korean	kor
Portuguese, Brazilian	ptb
Russian	rus
Spanish	esp

AIX

Linux

Windows

Requirements and limitations for IBM Spectrum Protect client management services

IBM Spectrum Protect™ client management services is a component that you install on backup-archive clients to collect diagnostic information such as client log files. Before you install the client management service on your system, you must understand the requirements and limitations.

In the documentation for the client management service, *client system* is the system where the backup-archive client is installed.

Diagnostic information can be collected only from Linux and Windows clients, but administrators can view the diagnostic information in the Operations Center on AIX®, Linux, or Windows operating systems.

Requirements for the client management service

Verify the following requirements before you install the client management service:

- To remotely access the client, the Operations Center administrator must have system authority or one of the following client authority levels:
 - Policy authority
 - Client owner authority
 - Client node access authority
- Ensure that the client system meets the following requirements:
 - The client management service can be installed only on client systems that run on Linux or Windows operating systems:
 - Linux x86 64-bit operating systems that are supported for the backup-archive client.
 - Windows 32-bit and 64-bit operating systems that are supported for the backup-archive client.
 - Transport Layer Security (TLS) 1.2 must be installed for transmission of data between the client management service and Operations Center. Basic authentication is provided and data and authentication information are encrypted through the SSL channel. TLS 1.2 is automatically installed along with the necessary SSL certificates when you install the client management service.
- On Linux client systems, you must have root user authority to install the client management service.
- For client systems that can have multiple client nodes, such as Linux client systems, ensure that each node name is unique on the client system.

Tip: After you install the client management service, you do not have to install it again because the service can discover multiple client options files.

Limitations of the client management service

The client management service provides basic services for collecting diagnostic information from backup-archive clients. The following limitations exist for the client management service:

- You can install the client management service only on systems with backup-archive clients, including backup-archive clients that are installed on data mover nodes for IBM Spectrum Protect for Virtual Environments: Data Protection for VMware.
- You cannot install the client management service on other IBM Spectrum Protect client components or products that do not have backup-archive clients.
- If the backup-archive clients are protected by a firewall, ensure that the Operations Center can connect to the backup-archive clients through the firewall by using the configured port for the client management service. The default port is 9028, but it can be changed.
- The client management service scans all client log files to locate entries for the previous 72-hour period.
- The Diagnosis page in the Operations Center provides basic troubleshooting information for backup-archive clients. However, for some backup issues, you might have to access the client system and obtain further diagnostic information.
- If the combined size of the client error log files and schedule log files on a client system is more than 500 MB, delays can occur in sending log records to the Operations Center. You can control the size of the log files by enabling log file pruning or wrapping by specifying the `errorlogretention` or `errorlogmax` client option.
- If you use the same client node name to connect to multiple IBM Spectrum Protect servers that are installed on the same server, you can view log files for only one of the client nodes.

For updates about the client management service, including requirements, limitations, and documentation updates, see technote 1963610.

Related tasks:

Collecting diagnostic information with IBM Spectrum Protect client management services

AIX | Linux | Windows

Administrator IDs that the Operations Center requires

An administrator must have a valid ID and password on the hub server to log in to the Operations Center. An administrator ID is also assigned to the Operations Center so that the Operations Center can monitor servers.

The Operations Center requires the following IBM Spectrum Protect™ administrator IDs:

Administrator IDs that are registered on the hub server

Any administrator ID that is registered on the hub server can be used to log in to the Operations Center. The authority level of the ID determines which tasks can be completed. You can create new administrator IDs by using the REGISTER ADMIN command.

Restriction: To use an administrator ID in a multiple-server configuration, the ID must be registered on the hub and spoke servers with the same password and authority level.

To manage authentication for these servers, consider using one of the following methods:

- A Lightweight Directory Access Protocol (LDAP) server
- The enterprise configuration functions to automatically distribute changes to the administrator definitions.

Monitoring administrator ID

When you initially configure the hub server, an administrator ID named IBM-OC-*server_name* is registered with system authority on the hub server and is associated with the initial password that you specify. This ID, which is sometimes called the *monitoring administrator*, is intended for use only by the Operations Center.

Do not delete, lock, or modify this ID. The same administrator ID with the same password is registered on the spoke servers that you add. The password is automatically changed on the hub and spoke servers every 90 days. You do not need to use or manage this password.

Restriction: The Operations Center maintains the monitoring administrator ID and password on spoke servers unless you use an enterprise configuration to manage these credentials. For more information about using an enterprise configuration to manage the credentials, see Tips for designing the hub and spoke server configuration.

Related reference:

REGISTER ADMIN (Register an administrator ID)

AIX

Linux

Windows

IBM Installation Manager

The Operations Center uses IBM® Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install the Operations Center. It must remain installed on the system so that the Operations Center can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The Operations Center offering contains all of the media that IBM Installation Manager requires to install the Operations Center.

Package

The group of software components that are required to install an offering.

The Operations Center package contains the following components:

- IBM Installation Manager installation program
- Operations Center offering

Package group

A set of packages that share a common parent directory.

Repository

A remote or local storage area for data and other application resources.

The Operations Center package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of the Operations Center.

AIX

Linux

Windows

Installation checklist

Before you install the Operations Center, you must verify certain information, such as the installation credentials, and you must determine the input to provide to IBM® Installation Manager for the installation.

The following checklist highlights the information that you must verify or determine before you install the Operations Center, and Table 1 describes the details of this information:

- Verify the host name for the computer where the Operations Center is to be installed.
- Verify the installation credentials.
- Determine the Operations Center installation directory, if you do not want to accept the default path.
- Determine the IBM Installation Manager installation directory, if you do not want to accept the default path.
- Determine the port number to be used by the Operations Center web server, if you do not want to accept the default port number.
- Determine the password for secure communications.

Table 1. Information to verify or determine before you install the Operations Center

Information	Details
Host name for the computer where the Operations Center is to be installed.	<p>The host name must meet the following criteria:</p> <ul style="list-style-type: none"> • It must not contain double-byte character set (DBCS) characters or the underscore character (_). • Although the host name can contain the hyphen character (-), it cannot have a hyphen as the last character in the name.
Installation credentials	<p>To install the Operations Center, you must use the following user account:</p> <ul style="list-style-type: none"> • AIX Linux The root user • Windows Administrator
Operations Center installation directory	<p>The Operations Center is installed in the ui subdirectory of the installation directory.</p> <p>The following path is the default path for the Operations Center installation directory:</p> <ul style="list-style-type: none"> • AIX Linux /opt/tivoli/tsm For example, if you use this default path, the Operations Center is installed in the following directory: <code>/opt/tivoli/tsm/ui</code> • Windows c:\Program Files\Tivoli\TSM For example, if you use this default path, the Operations Center is installed in the following directory: <code>c:\Program Files\Tivoli\TSM\ui</code> <p>The installation directory path must meet the following criteria:</p> <ul style="list-style-type: none"> • The path must contain no more than 128 characters. • The path must include only ASCII characters. • The path cannot include non-displayable control characters. • The path cannot include any of the following characters: <code>% < > ' " \$ & ; *</code>

Information	Details
IBM Installation Manager installation directory	<p>The following path is the default path for the IBM Installation Manager installation directory:</p> <ul style="list-style-type: none"> • AIX Linux /opt/IBM/InstallationManager • Windows C:\Program Files\IBM\Installation Manager
Port number that is used by the Operations Center web server.	<p>The value for the secure (https) port number must meet the following criteria:</p> <ul style="list-style-type: none"> • The number must be an integer in the range 1024 - 65535. • The number cannot be in use or allocated to other programs. <p>If you do not specify a port number, the default value is 11090.</p> <p>Tip: If you later do not remember the port number that you specified, refer to the following file, where <i>installation_dir</i> represents the directory where the Operations Center is installed:</p> <ul style="list-style-type: none"> • AIX Linux <i>installation_dir</i>/ui/Liberty/usr/servers/guiServer/bootstrap.properties • Windows <i>installation_dir</i>\ui\Liberty\usr\servers\guiServer\bootstrap.properties <p>The bootstrap.properties file contains the IBM Spectrum Protect™ server connection information.</p>
Password for secure communications	<p>The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers.</p> <p>The Operations Center requires secure communication between the server and the Operations Center. To secure communication, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.</p> <p>The truststore file of the Operations Center contains the certificate that the Operations Center uses for HTTPS communication with web browsers. During installation of the Operations Center, you create a password for the truststore file. When you set up secure communication between the Operations Center and the hub server, you must use the same password to add the certificate of the hub server to the truststore file.</p> <p>The password for the truststore file must meet the following criteria:</p> <ul style="list-style-type: none"> • The password must contain a minimum of 6 characters and a maximum of 64 characters. • The password must contain at least the following characters: <ul style="list-style-type: none"> ◦ One uppercase letter (A – Z) ◦ One lowercase letter (a – z) ◦ One digit (0 – 9) ◦ Two of the non-alphanumeric characters that are listed in the following series: <pre> ~ @ # \$ % ^ & * _ - + = ` () { } [] : ; < > , . ? / </pre>

Related tasks:

Configuring for secure communication

Resetting the password for the Operations Center truststore file

AIX | **Linux** | **Windows**

Installing the Operations Center

You can install the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

Before you begin

You cannot configure the Operations Center until you install, configure, and start the IBM Spectrum Protect™ server. Therefore, before you install the Operations Center, install the appropriate server package, according to the server version requirements in Hub and spoke server requirements.

You can install the Operations Center on a computer with the IBM Spectrum Protect server or on a separate computer.

- Obtaining the Operations Center installation package
You can obtain the installation package from an IBM® download site such as IBM Passport Advantage® or IBM Fix Central.
- Installing the Operations Center by using a graphical wizard
You can install or update the Operations Center by using the graphical wizard of IBM Installation Manager.
- Installing the Operations Center in console mode
You can install or update the Operations Center by using the command line in console mode.
- Installing the Operations Center in silent mode
You can install or upgrade the Operations Center in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

AIX Linux Windows

Obtaining the Operations Center installation package

You can obtain the installation package from an IBM® download site such as IBM Passport Advantage® or IBM Fix Central.

About this task

After you obtain the package from an IBM download site, you must extract the installation files.

Procedure

Complete the following steps to extract the Operations Center installation files. In the following steps, replace *version_number* with the version of Operations Center that you are installing.

AIX On AIX® systems:

- a. Download the following package file to the directory of your choice:

```
version_number.000  
-IBM-SPOC-AIX.bin
```

- b. Ensure that you have executable permission for the package file.
If necessary, change the file permissions by issuing the following command:

```
chmod a+x version_number.000-IBM-SPOC-AIX.bin
```

- c. Issue the following command to extract the installation files:

```
./version_number.000-IBM-SPOC-AIX.bin
```

The self-extracting package file is extracted to the directory.

Linux On Linux systems:

- a. Download one of the following package files to the directory of your choice:

- o *version_number.000-IBM-SPOC-LinuxS390.bin*
- o *version_number.000-IBM-SPOC-Linuxx86_64.bin*

- b. Ensure that you have executable permission for the package file.
If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

- c. Issue the following command to extract the installation files:

```
./package_name.bin
```

The self-extracting package file is extracted to the directory.

Windows On Windows systems:

- a. Download the following package file to the directory of your choice:

`version_number.000-IBM-SPOC-WindowsX64.exe`

b. In Windows Explorer, double-click the file name to extract the installation files.

The self-extracting package file is extracted to the directory.

AIX

Linux

Windows

Installing the Operations Center by using a graphical wizard

You can install or update the Operations Center by using the graphical wizard of IBM® Installation Manager.

AIX

Before you begin

If the following RPM files are not installed on the computer, install them. For instructions, see [Installing RPM files for the graphical wizard](#).

- `atk-1.12.3-2.aix5.2.ppc.rpm`
- `cairo-1.8.8-1.aix5.2.ppc.rpm`
- `expat-2.0.1-1.aix5.2.ppc.rpm`
- `fontconfig-2.4.2-1.aix5.2.ppc.rpm`
- `freetype2-2.3.9-1.aix5.2.ppc.rpm`
- `gettext-0.10.40-6.aix5.1.ppc.rpm`
- `glib2-2.12.4-2.aix5.2.ppc.rpm`
- `gtk2-2.10.6-4.aix5.2.ppc.rpm`
- `libjpeg-6b-6.aix5.1.ppc.rpm`
- `libpng-1.2.32-2.aix5.2.ppc.rpm`
- `libtiff-3.8.2-1.aix5.2.ppc.rpm`
- `pango-1.14.5-4.aix5.2.ppc.rpm`
- `pixman-0.12.0-3.aix5.2.ppc.rpm`
- `xcursor-1.1.7-3.aix5.2.ppc.rpm`
- `xft-2.1.6-5.aix5.1.ppc.rpm`
- `xrender-0.9.1-3.aix5.2.ppc.rpm`
- `zlib-1.2.3-3.aix5.1.ppc.rpm`

Procedure

1. From the directory where the Operations Center installation package file is extracted, issue the following command:

- `AIX Linux ./install.sh`
- `Windows install.bat`

2. Follow the wizard instructions to install the IBM Installation Manager and Operations Center packages.

AIX The following message might be displayed, and the installation wizard might be slow, if your locale uses UTF-8 encoding:

```
Cannot create font set
```

If the message is displayed, take one of the following actions:

- Change to a locale that does not use UTF-8 encoding. For language-option values that do not use UTF-8 encoding, see [Language requirements](#).
- Install the Operations Center by using the command line in console mode.
- Install the Operations Center in silent mode.

What to do next

See [Configuring the Operations Center](#).

- **AIX** [Installing RPM files for the graphical wizard](#)

Before you can use the graphical wizard of IBM Installation Manager to install the Operations Center, certain RPM files must be installed.

AIX

Linux

Windows

Installing the Operations Center in console mode

You can install or update the Operations Center by using the command line in console mode.

Procedure

1. From the directory where the installation package file is extracted, run the following program:

```
AIX | Linux  
./install.sh -c
```

```
Windows  
install.bat -c
```

2. Follow the console instructions to install the Installation Manager and Operations Center packages.

What to do next

See Configuring the Operations Center.

```
AIX | Linux | Windows
```

Installing the Operations Center in silent mode

You can install or upgrade the Operations Center in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the input directory where the installation package is extracted:

```
install_response_sample.xml  
    Use this file to install the Operations Center.  
update_response_sample.xml  
    Use this file to upgrade the Operations Center.
```

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
Tip: To generate a response file as part of a console-mode installation, complete the selection of the console-mode installation options. Then, in the Summary panel, enter **G** to generate the response file according to the previously selected options.
2. Create a password for the Operations Center truststore in the response file.
If you are using the `install_response_sample.xml` file, add the password in the following line of the file, where `mypassword` represents the password:

```
<variable name='ssl.password' value='mypassword' />
```

For more information about this password, see Installation checklist.

Tip: To upgrade the Operations Center, the truststore password is not required if you are using the `update_response_sample.xml` file.

3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value `response_file` represents the response file path and file name:

```
AIX | Linux  
./install.sh -s -input response_file -acceptLicense
```

o **Windows**

```
install.bat -s -input response_file -acceptLicense
```

What to do next

See Configuring the Operations Center.

AIX

Linux

Windows

Upgrading the Operations Center

You can upgrade the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

Before you begin

Before you upgrade the Operations Center, review the system requirements and the installation checklist. The new version of the Operations Center might have more or different requirements and considerations than the version you are currently using.

About this task

The instructions for upgrading the Operations Center are the same as the instructions for installing the Operations Center, with the following exceptions:

- You use the Update function of IBM® Installation Manager rather than the Install function.
Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.
- If you are upgrading the Operations Center in silent mode, you can skip the step of creating a password for the truststore file.

AIX

Linux

Windows

Getting started with the Operations Center

Before you can use the Operations Center to manage your storage environment, you must configure it.


About this task

After you install the Operations Center, complete the following basic configuration steps:

1. Designate the hub server.
2. Add any spoke servers.
3. Optionally, configure email alerts on the hub and spoke servers.

Figure 1 illustrates an Operations Center configuration.

Figure 1. Example of an Operations Center configuration with the hub and spoke servers



- **Configuring the Operations Center**
When you open the Operations Center for the first time, you must configure it to manage your storage environment. You must associate the Operations Center with the IBM Spectrum Protect server that is designated as the hub server. You can then connect additional IBM Spectrum Protect servers as spoke servers.
- **Configuring for secure communication**
The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers. The Transport Layer Security (TLS) protocol secures communications between the Operations Center and the hub server, and between the hub server and associated spoke servers.
- **Starting and stopping the web server**
The web server of the Operations Center runs as a service and starts automatically. You might need to stop and start the web server, for example, to make configuration changes.
- **Opening the Operations Center**
The Overview page is the default initial view in the Operations Center. However, in your web browser, you can bookmark the page that you want to open when you log in to the Operations Center.
- **Collecting diagnostic information with IBM Spectrum Protect client management services**
The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

AIX Linux Windows

Configuring the Operations Center

When you open the Operations Center for the first time, you must configure it to manage your storage environment. You must associate the Operations Center with the IBM Spectrum Protect™ server that is designated as the hub server. You can then connect additional IBM Spectrum Protect servers as spoke servers.

- **Designating the hub server**
When you connect to the Operations Center for the first time, you must designate which IBM Spectrum Protect server is the hub server.
- **Adding a spoke server**
After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.
- **Sending email alerts to administrators**
An alert is a notification of a relevant problem on the IBM Spectrum Protect server and is triggered by a server message. Alerts can be shown in the Operations Center and can be sent from the server to administrators by email.
- **Adding customized text to the login screen**
You can add customized text, such as your organization's Terms of Use of the software, to the login screen of the Operations Center so that users of the Operations Center see the text before they enter their user name and password.
- **Enabling REST services**
Applications that use Representational State Transfer (REST) services can query and manage the storage environment by connecting to the Operations Center.

Designating the hub server

When you connect to the Operations Center for the first time, you must designate which IBM Spectrum Protect™ server is the hub server.

Before you begin

The Operations Center requires secure communication between the hub server and the Operations Center. To secure communication, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center. For more information, see [Securing communication between the Operations Center and the hub server](#).

Procedure

In a web browser, enter the following address, where *hostname* represents the name of the computer where the Operations Center is installed, and *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer:

```
https://hostname:secure_port/oc
```

Tips:

- The URL is case-sensitive. For example, ensure that you type "oc" in lowercase as indicated.
- For more information about the port number, see the Installation checklist.
- If you are connecting to the Operations Center for the first time, you must provide the following information:
 - Connection information for the server that you want to designate as a hub server
 - Login credentials for an administrator ID that is defined for that server
- If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a hub server.

What to do next

If you have multiple IBM Spectrum Protect servers in your environment, add the other servers as spoke servers to the hub server.

Attention: Do not change the name of a server after it is configured as a hub or spoke server.

Related concepts:

[Hub and spoke server requirements](#)

[Administrator IDs that the Operations Center requires](#)

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of the hub server.

Procedure

1. In the Operations Center menu bar, click Servers. The Servers page opens.

In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke server.

2. Complete one of the following steps:
 - Click the server to highlight it, and in the table menu bar, click Monitor Spoke.
 - If the server that you want to add is not shown in the table, and secure SSL/TLS communication is not required, click + Spoke in the table menu bar.

3. Provide the necessary information, and complete the steps in the spoke configuration wizard.
Tip: If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Related reference:

DEFINE SERVER (Define a server for server-to-server communications)

AIX

Linux

Windows

Sending email alerts to administrators

An alert is a notification of a relevant problem on the IBM Spectrum Protect™ server and is triggered by a server message. Alerts can be shown in the Operations Center and can be sent from the server to administrators by email.

Before you begin

Before you configure email notification for administrators about alerts, ensure that the following requirements are met:

- An SMTP server is required to send and receive alerts by email, and the server that sends the alerts by email must have access to the SMTP server.
Tip: If the Operations Center is installed on a separate computer, that computer does not need access to the SMTP server.
- An administrator must have system privilege to configure email notification.

About this task

An email notification is sent only for the first occurrence of an alert. Also, if an alert is generated before you configure email notification, no email notification is sent for that alert.

You can configure email notification in the following ways:

- Send notification for individual alerts
- Send alert summaries

An alert summary contains information about current alerts. The summary includes the total number of alerts, the total number of active and inactive alerts, the oldest alert, the newest alert, and the most frequently occurring alert.

You can specify a maximum of three administrators to receive alert summaries by email. Alert summaries are sent approximately every hour.

Procedure

To configure email notification for administrators about alerts, complete the following steps on each hub and spoke server from which you want to receive email alerts:

1. To verify that alert monitoring is turned on, issue the following command:

```
QUERY MONITORSETTINGS
```

2. If the command output indicates that alert monitoring is turned off, issue the following command. Otherwise, proceed to the next step.

```
SET ALERTMONITOR ON
```

3. To enable the sending of email notification, issue the following command:

```
SET ALERTEMAIL ON
```

4. To define the SMTP server that is used to send email notification, issue the following command:

```
SET ALERTEMAILSMTPHOST host_name
```

5. To specify the port number for the SMTP server, issue the following command:

```
SET ALERTEMAILSMTPPORT port_number
```

The default port number is 25.

6. To specify the email address of the sender of the alerts, issue the following command:

```
SET ALERTEMAILFROMADDR email_address
```

7. For each administrator ID that must receive email notification, issue one of the following commands to activate email notification and to specify the email address:

```
REGISTER ADMIN admin_name ALERT=YES EMAILADDRESS=email_address
```

```
UPDATE ADMIN admin_name ALERT=YES EMAILADDRESS=email_address
```

8. Choose either, or both, of the following options, and specify the administrator IDs to receive email notification:

- o Send notification for individual alerts

To specify or update the administrator IDs to receive email notification for an individual alert, issue one of the following commands:

```
DEFINE ALERTTRIGGER message_number Admin=admin_name1,admin_name2
```

```
UPDATE ALERTTRIGGER message_number ADDadmin=admin_name3 DELadmin=admin_name1
```

Tip: From the Configure Alerts page of the Operations Center, you can select the administrators who will receive email notification.

- o Send alert summaries

To specify or update the administrator IDs to receive alert summaries by email, issue the following command:

```
SET ALERTSUMMARYTOADMINS admin_name1,admin_name2,admin_name3
```

If you want to receive alert summaries but do not want to receive notification about individual alerts, complete the following steps:

- a. Suspend notification about individual alerts, as described in Suspending email alerts temporarily.
- b. Ensure that the respective administrator ID is listed in the following command:

```
SET ALERTSUMMARYTOADMINS admin_name1,admin_name2,admin_name3
```

Sending email alerts to multiple administrators

The following example illustrates the commands that cause any alerts for message ANR1075E to be sent in an email to the administrators *myadmin*, *djadmin*, and *csadmin*:

```
SET ALERTMONITOR ON
SET ALERTEMAIL ON
SET ALERTEMAILSMTPHOST mymailserver.domain.com
SET ALERTEMAILSMTPPORT 450
SET ALERTEMAILFROMADDR srvadmin@mydomain.com
UPDATE ADMIN myadmin ALERT=YES EMAILADDRESS=myaddr@anycompany.com
UPDATE ADMIN djadmin ALERT=YES EMAILADDRESS=djaddr@anycompany.com
UPDATE ADMIN csadmin ALERT=YES EMAILADDRESS=csaddr@anycompany.com
DEFINE ALERTTRIGGER anr0175e ADMIN=myadmin,djadmin,csadmin
```

- Suspending email alerts temporarily

In certain situations, you might want to suspend email alerts temporarily. For example, you might want to receive alert summaries but suspend notification about individual alerts, or you might want to suspend email alerts when an administrator is on vacation.

Related reference:

DEFINE ALERTTRIGGER (Define an alert trigger)

QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)

REGISTER ADMIN (Register an administrator ID)

SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)

SET ALERTEMAILFROMADDR (Set the email address of the sender)

SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)

SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)

SET ALERTMONITOR (Set the alert monitor to on or off)

SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)

UPDATE ADMIN (Update an administrator)

UPDATE ALERTTRIGGER (Update a defined alert trigger)

AIX

Linux

Windows

Adding customized text to the login screen

You can add customized text, such as your organization's Terms of Use of the software, to the login screen of the Operations Center so that users of the Operations Center see the text before they enter their user name and password.

Procedure

To add customized text to the login screen, complete the following steps:

1. On the computer where the Operations Center is installed, go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

AIX | **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer`

Windows `installation_dir\ui\Liberty\usr\servers\guiServer`

2. In the directory, create a file that is named `loginText.html` that contains the text that you want to add to the login screen. Any special, non-ASCII text must be UTF-8 encoded.

Tip: You can format the text by adding HTML tags.

3. Review the added text on the login screen of the Operations Center.

To open the Operations Center, enter the following address in a web browser, where *hostname* represents the name of the computer where the Operations Center is installed, and *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer:

`https://hostname:secure_port/oc`

Enabling REST services

Applications that use Representational State Transfer (REST) services can query and manage the storage environment by connecting to the Operations Center.

About this task

Enable this feature to allow REST services to interact with hub and spoke servers by sending calls to the following address:


`https://oc_host_name:port/oc/api`

where *oc_host_name* is the network name or IP address of the Operations Center host system and *port* is the Operations Center port number. The default port number is 11090.

For information about the REST services that are available for the Operations Center, see Technote <http://www-01.ibm.com/support/docview.wss?uid=swg21997347>, or issue the following REST call:

`https://oc_host_name:port/oc/api/help`

Procedure

1. On the Operations Center menu bar, hover over the settings icon  and click Settings.
2. On the General page, select the Enable administrative REST API check box.
3. Click Save.

AIX | **Linux** | **Windows**

Configuring for secure communication

The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers. The Transport Layer Security (TLS) protocol secures communications between the Operations Center and the hub server, and between the hub server and associated spoke servers.

About this task

TLS 1.2 is required for secure communication between the IBM Spectrum Protect™ server and the Operations Center, and between the hub server and spoke servers.

- Securing communication between the Operations Center and the hub server
To secure communications between the Operations Center and the hub server, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.
- Securing communication between the hub server and a spoke server
To secure communications between the hub server and a spoke server by using the Transport Layer Security (TLS) protocol, you must define the certificate of the spoke server to the hub server, and the certificate of the hub server to the spoke server. You must also configure the Operations Center to monitor the spoke server.
- Resetting the password for the Operations Center truststore file
To set up secure communication between the Operations Center and the hub server, you must know the password for the truststore file of the Operations Center. You create this password during the installation of the Operations Center. If you do not know the password, you can reset it.

AIX

Linux

Windows

Securing communication between the Operations Center and the hub server

To secure communications between the Operations Center and the hub server, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.

Before you begin

The truststore file of the Operations Center is a container for certificates that the Operations Center can access. The truststore file contains the certificate that the Operations Center uses for HTTPS communication with web browsers.

During the installation of the Operations Center, you create a password for the truststore file. To secure communication between the Operations Center and the hub server, you must use the same password to add the certificate of the hub server to the truststore file. If you do not remember this password, you can reset it. See [Resetting the password for the Operations Center truststore file](#).

Procedure

1. Specify the cert256.arm certificate as the default certificate in the key database file of the hub server.

To specify cert256.arm as the default certificate, complete the following steps:

- a. Issue the following command from the hub server instance directory:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed
-label "TSM Server SelfSigned SHA Key"
```

- b. Restart the hub server so that it can receive the changes to the key database file.

2. To verify that the cert256.arm certificate is set as the default certificate in the key database file of the hub server, issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

3. Stop the Operations Center web server.
4. Go to the command line of the operating system on which the Operations Center is installed.
5. Add the certificate to the truststore file of the Operations Center by using the iKeycmd utility or the iKeyman utility.

The iKeycmd utility is a command line interface, and the iKeyman utility is the IBM® Key Management graphical user interface.

AIX

Linux

The iKeycmd and the iKeyman utilities must be run as the root user.

Windows

The iKeycmd and the iKeyman utilities must be run by an administrator account.

To add the TLS certificate by using the command line interface, complete the following steps:

- a. Go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

- **AIX** **Linux** *installation_dir*/ui/jre/bin
- **Windows** *installation_dir*\ui\jre\bin

- b. Issue the ikeycmd command to add the cert256.arm certificate as the default certificate in the key database file of the hub server:

```
ikeycmd -cert -add
-db /installation_dir/Liberty/usr/servers/guiServer/gui-truststore.jks
-file /fvt/comfrey/srv/cert256.arm
-label 'label description'
-pw 'password' -type jks -format ascii -trust enable
```

where:

installation_dir

The directory in which the Operations Center is installed.

label description

The description that you assign to the label.

password

The password that you created when you installed the Operations Center. To reset the password, uninstall the Operations Center, delete the .jks file, and reinstall the Operations Center.

To add the certificate by using the IBM Key Management window, complete the following steps:

- a. Go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

- | | |
|-----|-------|
| AIX | Linux |
|-----|-------|

installation_dir/ui/jre/bin
- | | |
|---------|--|
| Windows | |
|---------|--|

installation_dir\ui\jre\bin

- b. Open the IBM Key Management window by issuing the following command:

```
ikeyman
```

- c. Click Key Database File > Open.

- d. In the Open window, click Browse, and go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

- | | |
|-----|-------|
| AIX | Linux |
|-----|-------|

installation_dir/ui/Liberty/usr/servers/guiServer
- | | |
|---------|--|
| Windows | |
|---------|--|

installation_dir\ui\Liberty\usr\servers\guiServer

- e. In the guiServer directory, select the gui-truststore.jks file.

- f. Click Open, and click OK.

- g. Enter the password for the truststore file, and click OK.

- h. In the Key database content area of the IBM Key Management window, click the arrow, and select Signer Certificates from the list.

- i. Click Add.

- j. In the Open window, click Browse, and go to the hub server instance directory, as shown in the following example:

- | | |
|-----|-------|
| AIX | Linux |
|-----|-------|

 /opt/tivoli/tsm/server/bin
- | | |
|---------|--|
| Windows | |
|---------|--|

 c:\Program Files\Tivoli\TSM\server1

The directory contains the cert256.arm certificate.

If you cannot access the hub server instance directory from the Open window, complete the following steps:

- i. Use FTP or another file-transfer method to copy the cert256.arm files from the hub server to the following directory on the computer where the Operations Center is installed:

- | | |
|-----|-------|
| AIX | Linux |
|-----|-------|

installation_dir/ui/Liberty/usr/servers/guiServer
- | | |
|---------|--|
| Windows | |
|---------|--|

installation_dir\ui\Liberty\usr\servers\guiServer

- ii. In the Open window, go to the guiServer directory.

- k. Select the cert256.arm certificate as the certificate.

Tip: The certificate that you choose must be set as the default certificate in the key database file of the hub server. For more information, see step 1 and 2.

- l. Click Open, and click OK.

- m. Enter a label for the certificate. For example, enter the name of the hub server.

- n. Click OK. The SSL certificate of the hub server is added to the truststore file, and the label is displayed in the Key database content area of the IBM Key Management window.

- o. Close the IBM Key Management window.

6. Start the Operations Center web server.

7. When you connect to the Operations Center for the first time, you are prompted to identify the IP address or network name of the hub server, and the port number for communicating with the hub server. If the ADMINONCLIENTPORT server option is enabled for the IBM Spectrum Protect™ server, enter the port number that is specified by the TCPADMINPORT server

option. If the ADMINONCLIENTPORT server option is not enabled, enter the port number that is specified by the TCPPOINT server option.

If the Operations Center was previously configured, you can review the contents of the serverConnection.properties file to verify the connection information. The serverConnection.properties file is in the following directory on the computer where the Operations Center is installed:

- o

AIX	Linux
-----	-------

`installation_dir/ui/Liberty/usr/servers/guiServer`
- o

Windows

`installation_dir\ui\Liberty\usr\servers\guiServer`

What to do next

To set up TLS communication between the hub server and a spoke server, see [Securing communication between the hub server and a spoke server](#).

Related reference:

QUERY OPTION (Query server options)

AIX	Linux	Windows
-----	-------	---------

Securing communication between the hub server and a spoke server

To secure communications between the hub server and a spoke server by using the Transport Layer Security (TLS) protocol, you must define the certificate of the spoke server to the hub server, and the certificate of the hub server to the spoke server. You must also configure the Operations Center to monitor the spoke server.

About this task

The hub server receives status and alert information from the spoke server and shows this information in the Operations Center. To receive the status and alert information from the spoke server, the certificate of the spoke server must be added to the truststore file of the hub server. You must also configure the Operations Center to monitor the spoke server.

To enable other functions of the Operations Center, such as the automatic deployment of client updates, the certificate of the hub server must be added to the truststore file of the spoke server.

Procedure

1. Complete the following steps to define the certificate of the spoke server to the hub server:
 - a. On the spoke server, change to the directory of the spoke server instance.
 - b. Specify the required cert256.arm certificate as the default certificate in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed  
-label "TSM Server SelfSigned SHA Key"
```

- c. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- d. Securely transfer the cert256.arm file of the spoke server to the hub server.
 - e. On the hub server, change to the directory of the hub server instance.
 - f. Define the spoke server certificate to the hub server. Issue the following command from the hub server instance directory, where *spoke_servername* is the name of the spoke server, and *spoke_cert256.arm* is the file name of the spoke server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label spoke_servername -file spoke_cert256.arm
```

2. Complete the following steps to define the certificate of the hub server to the spoke server:
 - a. On the hub server, change to the directory of the hub server instance.
 - b. Specify the required cert256.arm certificate as the default certificate in the key database file of the hub server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed  
-label "TSM Server SelfSigned SHA Key"
```

- c. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- d. Securely transfer the cert256.arm file of the hub server to the spoke server.
- e. On the spoke server, change to the directory of the spoke server instance.
- f. Define the hub server certificate to the spoke server. Issue the following command from the spoke server instance directory, where *hub_servername* is the name of the hub server, and *hub_cert256.arm* is the file name of the hub server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label hub_servername -file hub_cert256.arm
```

3. Restart the hub server and the spoke server.
4. Complete the following steps to define the spoke server to the hub server, and the hub server to the spoke server.
 - a. Issue the following commands on both the hub server and the spoke server:

```
SET SERVERPASSWORD server_password  
SET SERVERHLADDRESS ip_address  
SET SERVERLLADDRESS tcp_port
```

- b. On the hub server, issue the DEFINE SERVER command, according to the following example:

```
DEFINE SERVER spoke_servername HLA=spoke_address  
LLA=spoke_SSLTCPADMINPort SERVERPA=spoke_serverpassword
```

- c. On the spoke server, issue the DEFINE SERVER command, according to the following example:

```
DEFINE SERVER hub_servername HLA=hub_address  
LLA=hub_SSLTCPADMINPort SERVERPA=hub_serverpassword
```

Tip: By default, server communication is encrypted except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure. To encrypt all communication with the specified server, even when the server is sending and receiving object data, specify the SSL=YES parameter on the DEFINE SERVER command.

5. Complete the following steps to configure the Operations Center to monitor the spoke server:
 - a. On the Operations Center menu bar, click Servers. The spoke server has a status of "Unmonitored." This status means that, although this server was defined to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke.
 - b. Click the spoke server to highlight the item, and click Monitor Spoke.

Related reference:

DEFINE SERVER (Define a server for server-to-server communications)

QUERY OPTION (Query server options)

[AIX](#) | [Linux](#) | [Windows](#)

Resetting the password for the Operations Center truststore file

To set up secure communication between the Operations Center and the hub server, you must know the password for the truststore file of the Operations Center. You create this password during the installation of the Operations Center. If you do not know the password, you can reset it.

About this task

To reset the password, you must create a new password, delete the truststore file of the Operations Center, and restart the Operations Center web server.

Attention: Complete these steps only if you do not know the truststore password. Do not complete these steps if you know the truststore password and want only to change the password. To reset the password, you must delete the truststore file, which deletes all certificates that are already stored in the truststore file. If you know the truststore password, you can change it by using the iKeycmd or the iKeyman utility.

Procedure

1. Stop the Operations Center web server.
2. Go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:
 - o [AIX](#) | [Linux](#) | *installation_dir/ui/Liberty/usr/servers/guiServer*

- o **Windows** `installation_dir\ui\Liberty\usr\servers\guiServer`
3. Open the `bootstrap.properties` file, which contains the password for the truststore file. If the password is unencrypted, you can use it to open the truststore file without having to reset it.
- The following examples indicate the difference between an encrypted and an unencrypted password:

Encrypted password example

Encrypted passwords begin with the text string `{xor}`.

The following example shows the encrypted password as the value of the `tsm.truststore.pswd` parameter:

```
tsm.truststore.pswd={xor}MiYPPiwsKDAtoW==
```

Unencrypted password example

The following example shows the unencrypted password as the value of the `tsm.truststore.pswd` parameter:

```
tsm.truststore.pswd=J8b%^B
```

4. Reset the password by replacing the password in the `bootstrap.properties` file with a new password. You can replace the password with an encrypted or unencrypted password. Remember the unencrypted password for future use.

To create an encrypted password, complete the following steps:

- a. Create an unencrypted password.

The password for the truststore file must meet the following criteria:

- The password must contain a minimum of 6 characters and a maximum of 64 characters.
- The password must contain at least the following characters:
 - One uppercase letter (A – Z)
 - One lowercase letter (a – z)
 - One digit (0 – 9)
 - Two of the non-alphanumeric characters that are listed in the following series:

```
~ @ # $ % ^ & * _ - + = ` |
( ) { } [ ] : ; < > , . ? /
```

- b. From the command line of the operating system, go to the following directory:

- **AIX** | **Linux** `installation_dir/ui/Liberty/bin`
- **Windows** `installation_dir\ui\Liberty\bin`

- c. To encrypt the password, issue the following command, where `myPassword` represents the unencrypted password:

- **AIX** | **Linux** `securityUtility encode myPassword --encoding=aes`
- **Windows** `securityUtility.bat encode myPassword --encoding=aes`

Windows The following message might be shown:

```
! "java" is not recognized as an internal or external command,
operable program or batch file.
```

If this message is shown, complete the following steps:

- i. Issue the following command, where `installation_dir` represents the directory where the Operations Center is installed:

```
set JAVA_HOME="installation_dir\ui\jre"
```

- ii. Reissue the following command to encrypt the password:

```
securityUtility.bat encode myPassword --encoding=aes
```

5. Close the `bootstrap.properties` file.

6. Go to the following directory:

- o **AIX** | **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer`
- o **Windows** `installation_dir\ui\Liberty\usr\servers\guiServer`

7. Delete the `gui-truststore.jks` file, which is the truststore file of the Operations Center.

8. Start the Operations Center web server.

Results

A new truststore file is automatically created for the Operations Center, and the TLS certificate of the Operations Center is automatically included in the truststore file.

AIX | **Linux** | **Windows**

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might need to stop and start the web server, for example, to make configuration changes.

Procedure

Stop and start the web server.

- **AIX** From the `/installation_dir/ui/utils` directory, where `installation_dir` represents the directory where the Operations Center is installed, issue the following commands:

- To stop the server:

```
./stopserver.sh
```

- To start the server:

```
./startserver.sh
```

- **Linux** Issue the following commands:

- To stop the server:

```
service opscenter.rc stop
```

- To start the server:

```
service opscenter.rc start
```

- To restart the server:

```
service opscenter.rc restart
```

To determine whether the server is running, issue the following command:

```
service opscenter.rc status
```

- **Windows** From the Services window, stop or start the Operations Center service.

Opening the Operations Center

The Overview page is the default initial view in the Operations Center. However, in your web browser, you can bookmark the page that you want to open when you log in to the Operations Center.

Procedure

1. In a web browser, enter the following address, where `hostname` represents the name of the computer where the Operations Center is installed, and `secure_port` represents the port number that the Operations Center uses for HTTPS communication on that computer:

```
https://hostname:secure_port/oc
```

Tips:

- The URL is case-sensitive. For example, ensure that you type "oc" in lowercase as indicated.
- The default port number for HTTPS communication is 11090, but a different port number can be specified during Operations Center installation.

2. Log in, using an administrator ID that is registered on the hub server.

In the Overview page, you can view summary information for clients, services, servers, storage pools, and storage devices. You can view more details by clicking items or by using the Operations Center menu bar.

Monitoring from a mobile device: To remotely monitor the storage environment, you can view the Overview page of the Operations Center in the web browser of a mobile device. The Operations Center supports the Apple Safari web browser on the iPad. Other mobile devices can also be used.

AIX

Linux

Windows

Collecting diagnostic information with IBM Spectrum Protect client management services

The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

About this task

After you install the client management service, you can view the Diagnosis page in the Operations Center to obtain troubleshooting information for backup-archive clients.

Diagnostic information can be collected only from Linux and Windows clients, but administrators can view the diagnostic information in the Operations Center on AIX®, Linux, or Windows operating systems.

You can also install the client management service on data mover nodes for IBM Spectrum Protect™ for Virtual Environments: Data Protection for VMware to collect diagnostic information about the data movers.

Tip: In the documentation for the client management service, *client system* is the system where the backup-archive client is installed.

- Installing the client management service by using a graphical wizard
To collect diagnostic information about backup-archive clients such as client log files, you must install the client management service on the client systems that you manage.
- Installing the client management service in silent mode
You can install the client management service in silent mode. When you use silent mode, you provide the installation values in a response file and then run an installation command.
- Verifying that the client management service is installed correctly
Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.
- Configuring the Operations Center to use the client management service
If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.
- Starting and stopping the client management service
The client management service is automatically started after it is installed on the client system. You might need to stop and start the service in certain situations.
- Uninstalling the client management service
If you no longer have to collect client diagnostic information, you can uninstall the client management service from the client system.
- Configuring the client management service for custom client installations
The client management service uses information in the client configuration file (client-configuration.xml) to discover diagnostic information. If the client management service is unable to discover the location of log files, you must run the CmsConfig utility to add the location of the log files to the client-configuration.xml file.

AIX

Linux

Windows

Installing the client management service by using a graphical wizard

To collect diagnostic information about backup-archive clients such as client log files, you must install the client management service on the client systems that you manage.

Before you begin

Review Requirements and limitations for IBM Spectrum Protect client management services.

About this task

You must install the client management service on the same computer as the backup-archive client.

Procedure

1. Download the installation package for the client management service from an IBM® download site such as IBM Passport Advantage® or IBM Fix Central. Look for a file name that is similar to `<version>-IBM-SPCMS-<operating system>.bin`.

The following table shows the names of the installation packages.

Client operating system	Installation package name
Linux x86 64-bit	8.1.x.000-IBM-SPCMS-Linuxx64.bin
Windows 32-bit	8.1.x.000-IBM-SPCMS-Windows32.exe
Windows 64-bit	8.1.x.000-IBM-SPCMS-Windows64.exe

2. Create a directory on the client system that you want to manage, and copy the installation package there.
3. Extract the contents of the installation package file.

- o On Linux client systems, complete the following steps:
 - a. Change the file to an executable file by issuing the following command:

```
chmod +x 8.1.x.000-IBM-SPCMS-Linuxx64.bin
```

- b. Issue the following command:

```
./8.1.x.000-IBM-SPCMS-Linuxx64.bin
```

- o On Windows client systems, double-click the installation package name in Windows Explorer.

Tip: If you previously installed and uninstalled the package, select All when prompted to replace the existing installation files.

4. Run the installation batch file from the directory where you extracted the installation files and associated files. This is the directory that you created in step 2.

- o On Linux client systems, issue the following command:

```
./install.sh
```

- o On Windows client systems, double-click install.bat.

5. To install the client management service, follow the instructions in the IBM Installation Manager wizard.

If IBM Installation Manager is not already installed on the client system, you must select both IBM Installation Manager and IBM Spectrum Protect Client Management Services.

Tip: You can accept the default locations for the shared resources directory and the installation directory for IBM Installation Manager.

What to do next

Follow the instructions in Verifying that the client management service is installed correctly.

AIX

Linux

Windows

Installing the client management service in silent mode

You can install the client management service in silent mode. When you use silent mode, you provide the installation values in a response file and then run an installation command.

Before you begin

Review Requirements and limitations for IBM Spectrum Protect client management services.

Extract the installation package by following the instructions in Installing the client management service by using a graphical wizard.

About this task

You must install the client management service on the same computer as the backup-archive client.

The input directory, which is in the directory where the installation package is extracted, contains the following sample response file:

install_response_sample.xml

You can use the sample file with the default values, or you can customize it.

Tip: If you want to customize the sample file, create a copy of the sample file, rename it, and edit the copy.

Procedure

1. Create a response file based on the sample file, or use the sample file, `install_response_sample.xml`. In either case, ensure that the response file specifies the port number for the client management service. The default port is 9028. For example:

```
<variable name='port' value='9028' />
```

2. Run the command to install the client management service and accept the license. From the directory where the installation package file is extracted, issue the following command, where `response_file` represents the response file path, including the file name:

On a Linux client system:

```
./install.sh -s -input response_file -acceptLicense
```

For example:

```
./install.sh -s -input /cms_install/input/install_response.xml -acceptLicense
```

On a Windows client system:

```
install.bat -s -input response_file -acceptLicense
```

For example:

```
install.bat -s -input c:\cms_install\input\install_response.xml -acceptLicense
```

What to do next

Follow the instructions in [Verifying that the client management service is installed correctly](#).

AIX | Linux | Windows

Verifying that the client management service is installed correctly

Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.

Procedure

On the client system, at the command line, run the following commands to view the configuration of the client management service:

- On Linux client systems, issue the following command:

```
client_install_dir/cms/bin/CmsConfig.sh list
```

where `client_install_dir` is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
/opt/tivoli/tsm/cms/bin/CmsConfig.sh list
```

The output is similar to the following text:

```
Listing CMS configuration
```

```
server1.example.com:1500 NO_SSL HOSTNAME
```

```
Capabilities: [LOG_QUERY]
```

```
Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys
```

```
Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
```

```
en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

```
Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

- On Windows client systems, issue the following command:

```
client_install_dir\cms\bin\CmsConfig.bat list
```

where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
C:\Program Files\Tivoli\TSM\cms\bin\CmsConfig.bat list
```

The output is similar to the following text:

```
Listing CMS configuration

server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt

  Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
  en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Log File: C:\Program Files\Tivoli\TSM\baclient\dmsched.log
  en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

If the client management service is correctly installed and configured, the output displays the location of the error log file. The output text is extracted from the following configuration file:

- On Linux client systems:

```
client_install_dir/cms/Liberty/usr/servers/cmsServer/client-configuration.xml
```

- On Windows client systems:

```
client_install_dir\cms\Liberty\usr\servers\cmsServer\client-configuration.xml
```

If the output does not contain any entries, you must configure the client-configuration.xml file. For instructions about how to configure this file, see [Configuring the client management service for custom client installations](#). You can use the `CmsConfig verify` command to verify that a node definition is correctly created in the client-configuration.xml file.

AIX | Linux | Windows

Configuring the Operations Center to use the client management service

If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Before you begin

Ensure that the client management service is installed and started on the client system.

Verify whether the default configuration is used. The default configuration is not used if either of the following conditions is met:

- The client management service does not use the default port number, 9028.
- The backup-archive client is not accessed by the same IP address as the client system where the backup-archive client is installed. For example, a different IP address might be used in the following situations:
 - The computer system has two network cards. The backup-archive client is configured to communicate on one network, while the client management service communicates on the other network.
 - The client system is configured with the Dynamic Host Configuration Protocol (DHCP). As a result, the client system is dynamically assigned an IP address, which is saved on the IBM Spectrum Protect™ server during the previous backup-archive client operation. When the client system is restarted, the client system might be assigned a different IP address. To ensure that the Operations Center can always find the client system, you specify a fully qualified domain name.

Procedure

To configure the Operations Center to use the client management service, complete the following steps:

1. On the Clients page of the Operations Center, select the client.
2. Click Details.
3. Click the Properties tab.
4. In the Remote diagnostics URL field in the General section, specify the URL for the client management service on the client system.

The address must start with `https`. The following table shows examples of the remote diagnostics URL.

Type of URL	Example
With DNS host name and default port, 9028	<code>https://server.example.com</code>
With DNS host name and non-default port	<code>https://server.example.com:1599</code>
With IP address and non-default port	<code>https://192.0.2.0:1599</code>

5. Click Save.

What to do next

You can access client diagnostic information such as client log files from the Diagnosis tab in the Operations Center.

AIX | Linux | Windows

Starting and stopping the client management service

The client management service is automatically started after it is installed on the client system. You might need to stop and start the service in certain situations.

Procedure

- To stop, start, or restart the client management service on Linux client systems, issue the following commands:
 - To stop the service:

```
service cms.rc stop
```
 - To start the service:

```
service cms.rc start
```
 - To restart the service:

```
service cms.rc restart
```
- On Windows client systems, open the Services window, and stop, start, or restart the IBM Spectrum Protect™ Client Management Services service.

AIX | Linux | Windows

Uninstalling the client management service

If you no longer have to collect client diagnostic information, you can uninstall the client management service from the client system.

About this task

You must use IBM® Installation Manager to uninstall the client management service. If you no longer plan to use IBM Installation Manager, you can also uninstall it.

Procedure

1. Uninstall the client management service from the client system:
 - a. Open IBM Installation Manager:

- On the Linux client system, in the directory where IBM Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command:

```
./IBMIM
```

- On the Windows client system, open IBM Installation Manager from the Start menu.
 - b. Click Uninstall.
 - c. Select IBM Spectrum Protect Client Management Services, and click Next.
 - d. Click Uninstall, and then click Finish.
 - e. Close the IBM Installation Manager window.
- 2. If you no longer require IBM Installation Manager, uninstall it from the client system:
 - a. Open the IBM Installation Manager uninstall wizard:
 - On the Linux client system, change to the IBM Installation Manager uninstallation directory (for example, /var/ibm/InstallationManager/uninstall), and issue the following command:

```
./uninstall
```
 - On the Windows client system, click Start > Control Panel. Then, click Uninstall a program > IBM Installation Manager > Uninstall.
 - b. In the IBM Installation Manager window, select IBM Installation Manager if it is not already selected, and click Next.
 - c. Click Uninstall, and click Finish.

AIX

Linux

Windows

Configuring the client management service for custom client installations

The client management service uses information in the client configuration file (client-configuration.xml) to discover diagnostic information. If the client management service is unable to discover the location of log files, you must run the CmsConfig utility to add the location of the log files to the client-configuration.xml file.

- CmsConfig utility
If you are not using the default client configuration, you can run the CmsConfig utility on the client system to discover and add the location of the client log files to the client-configuration.xml file. After you complete the configuration, the client management service can access the client log files and make them available for basic diagnostic functions in the Operations Center.

AIX

Linux

Troubleshooting the Operations Center installation

If a problem occurs with the Operations Center installation and you cannot solve it, you can consult the descriptions of known problems for a possible solution.

- **AIX** Graphical installation wizard cannot be started on an AIX system
You are installing the Operations Center on an AIX® system by using the graphical wizard, and the installation program does not start.
- **Linux** Chinese, Japanese, or Korean fonts are displayed incorrectly
Chinese, Japanese, or Korean fonts are displayed incorrectly in the Operations Center on Red Hat Enterprise Linux 5.

AIX

Graphical installation wizard cannot be started on an AIX system

You are installing the Operations Center on an AIX® system by using the graphical wizard, and the installation program does not start.

Solution

The RPM files that are listed in Installing the Operations Center by using a graphical wizard must be installed on the computer. Verify that the RPM files are installed.

Linux

Chinese, Japanese, or Korean fonts are displayed incorrectly

Chinese, Japanese, or Korean fonts are displayed incorrectly in the Operations Center on Red Hat Enterprise Linux 5.

Solution

Install the following font packages, which are available from Red Hat:

- fonts-chinese
- fonts-japanese
- fonts-korean

AIX | Linux | Windows

Uninstalling the Operations Center

You can uninstall the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

- Uninstalling the Operations Center by using a graphical wizard
You can uninstall the Operations Center by using the graphical wizard of IBM® Installation Manager.
- Uninstalling the Operations Center in console mode
To uninstall the Operations Center by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.
- Uninstalling the Operations Center in silent mode
To uninstall the Operations Center in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.

AIX | Linux | Windows

Uninstalling the Operations Center by using a graphical wizard

You can uninstall the Operations Center by using the graphical wizard of IBM® Installation Manager.

Procedure

1. Open IBM Installation Manager.

AIX | Linux In the directory where IBM Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command:

```
./IBMIM
```

Windows You can open IBM Installation Manager from the Start menu.

2. Click Uninstall.
3. Select the option for the Operations Center, and click Next.
4. Click Uninstall.
5. Click Finish.

AIX | Linux | Windows

Uninstalling the Operations Center in console mode

To uninstall the Operations Center by using the command line, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o AIX | Linux eclipse/tools
 - o Windows eclipse\tools

For example:

- o **AIX** **Linux** /opt/IBM/InstallationManager/eclipse/tools
 - o **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools
2. From the tools directory, issue the following command:
 - o **AIX** **Linux** ./imcl -c
 - o **Windows** imcl.exe -c
 3. To uninstall, enter 5.
 4. Choose to uninstall from the IBM Spectrum Protect™ package group.
 5. Enter N for Next.
 6. Choose to uninstall the Operations Center package.
 7. Enter N for Next.
 8. Enter U for Uninstall.
 9. Enter F for Finish.

AIX **Linux** **Windows**

Uninstalling the Operations Center in silent mode

To uninstall the Operations Center in silent mode, you must run the uninstallation program of IBM® Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the Operations Center server. IBM Spectrum Protect™ includes a sample response file, `uninstall_response_sample.xml`, in the input directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

To uninstall the Operations Center, leave `modify="false"` set for the Operations Center entry in the response file.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
 - o **AIX** **Linux** eclipse/tools
 - o **Windows** eclipse\tools

For example:

- o **AIX** **Linux** /opt/IBM/InstallationManager/eclipse/tools
 - o **Windows** C:\Program Files\IBM\Installation Manager\eclipse\tools
2. From the tools directory, issue the following command, where *response_file* represents the response file path, including the file name:
 - o **AIX** **Linux** ./imcl -input *response_file* -silent
 - o **Windows** imcl.exe -input *response_file* -silent
- The following command is an example:
- o **AIX** **Linux** ./imcl -input /tmp/input/uninstall_response.xml -silent
 - o **Windows** imcl.exe -input C:\tmp\input\uninstall_response.xml -silent

AIX **Linux** **Windows**

Rolling back to a previous version of the Operations Center

By default, IBM® Installation Manager saves earlier versions of a package to roll back to if you experience a problem with later versions of updates, fixes, or packages.

Before you begin

The rollback function is available only after the Operations Center is updated.

About this task

When IBM Installation Manager rolls back a package to a previous version, the current version of the package files is uninstalled, and an earlier version is reinstalled.

To roll back to a previous version, IBM Installation Manager must access files for that version. By default, these files are saved during each successive installation. Because the number of saved files increases with each installed version, you might want to delete these files from your system on a regular schedule. However, if you delete the files, you cannot roll back to a previous version.

To delete saved files or to update your preference for saving these files in future installations, complete the following steps:

1. In IBM Installation Manager, click File > Preferences.
2. On the Preferences page, click Files for Rollback, and specify your preference.

Procedure

To roll back to a previous version of the Operations Center, use the Roll Back function of IBM Installation Manager.

Configuring servers

To complete configuration tasks for the IBM Spectrum Protect™ server, review available documentation.

About this task

Tip: Beginning with IBM® Tivoli® Storage Manager Version 7.1.3, the *Administrator's Guide* topics are no longer available in PDF format. Instead, the documentation set is revised to help you complete specific tasks:

- To implement a new data protection solution, see IBM Spectrum Protect data protection solutions. The solution guides provide cookbook-style instructions to help you plan, implement, and manage a solution.
- Alternatively, you can use the IBM Spectrum Protect Blueprints. You can follow the Blueprint procedures to deploy a storage environment, and use the Blueprint scripts to streamline the installation and configuration process. The Blueprints provide the latest hardware and software requirements for small, medium, and large storage environments.
- To administer an *existing* solution, see the following table.

Action	Details	Documentation
Secure the server.	Secure the IBM Spectrum Protect server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.	Securing the IBM Spectrum Protect server
Learn about and configure policies for data retention.	IBM Spectrum Protect policies define the rules for managing your data.	To update policies, use the Operations Center. To learn more about policies and create policies, see Customizing policies.
Eliminate duplicate data.	Use data deduplication to eliminate redundant data in storage pools. Data deduplication reduces the storage that is required to retain the data. Only one instance of the data is retained in a deduplicated storage pool. With IBM Spectrum Protect V7.1.3 and later, you can use inline data deduplication.	To learn more about the differences between inline and post-process data deduplication and to configure the best practice solution for data deduplication, see Data deduplication options.

Action	Details	Documentation
Replicate data.	You can replicate client node data from a source replication server to a target replication server. If a disaster occurs and the source server is temporarily unavailable, client nodes can recover their data from the target replication server.	To implement a best practice solution that uses IBM Spectrum Protect replication and automatic failover, see Multisite disk solution. For general information about replication, including configuration steps, see Replicating client data to another server.
Monitor a storage solution.	Monitor the storage solution to identify existing and potential issues. In this way, you can troubleshoot problems and optimize system performance.	Monitoring storage solutions
Manage the database and recovery log.	The database and recovery log, or server inventory, store information about client data and are critical to the operation of the server.	<ul style="list-style-type: none"> • For general information about the database and recovery log, see: Managing the database and recovery log (V7.1.1). • To optimize index and table reorganization for the server database, and prevent and resolve issues that are related to database growth and performance issues, see technote 1683633.
Configure SSL for Lightweight Directory Access Protocol (LDAP).	You can configure SSL for LDAP directory servers and manage passwords and logon procedures.	For information about LDAP, see: <ul style="list-style-type: none"> • Authenticating users by using an LDAP server • Configuring SSL or TLS for LDAP directory servers (V7.1.1)
Protect the server and recover in a disaster.	Protect your system infrastructure and data so that you can recover from a disaster. Use the tools and procedures that IBM Spectrum Protect provides to help you create a disaster plan.	For information about protecting and recovering the server and data, see: <ul style="list-style-type: none"> • Protecting the database and infrastructure setup files (V7.1.1) • Using disaster recovery manager for tape environments (V7.1.1) • Repairing and recovering data
Protect clients.	The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.	Configuring clients for applications, virtual machines, and systems
Select and configure storage.	Select storage based on your business needs and then complete the tasks for configuration.	Configuring storage
Protect NAS file servers.	You can plan, configure, and manage a backup environment that protects your network-attached storage (NAS) file server.	Protecting NAS file servers

Action	Details	Documentation
Configure a clustered environment.	Configure a clustered environment on AIX®, Linux, or Windows operating systems to ensure higher server availability and minimized downtime.	Configuring clustered environments
Configure virtual tape libraries.	A virtual tape library (VTL) does not use physical tape media. When you implement VTL storage, you can exceed the capacity of a physical tape library. The ability to define many volumes and drives can provide greater flexibility for the storage environment.	Configuring virtual tape libraries
Protect data with the NetApp SnapLock licensed feature.	You can use the NetApp SnapLock licensed feature to meet strict regulatory requirements for archived data.	Data protection by using the NetApp SnapLock licensed feature
Manage operations.	Manage server and client operations to prevent potential issues and improve performance.	Managing operations

- Securing the IBM Spectrum Protect server
Secure the IBM Spectrum Protect server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.
- Replicating client data to another server
Replicating client data from a source server to another server helps to ensure that backed-up client data is available for recovery if the source server is damaged. Replication incrementally copies data from the source server to the target server to provide failover and failback capability.
- Configuring clustered environments
You can configure the IBM Spectrum Protect server for clustering on AIX, Linux, or Windows systems.

Securing the IBM Spectrum Protect server

Secure the IBM Spectrum Protect™ server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.

- Security concepts
You can protect IBM Spectrum Protect from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.
- Managing administrators
An administrator who has system authority can complete any task with the IBM Spectrum Protect server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.
- Changing password requirements
You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect.
- Securing IBM Spectrum Protect on the system
Protect the system where the IBM Spectrum Protect server runs to prevent unauthorized access.
- Protecting the storage environment against ransomware
Storage environments that are connected to the internet can be the target of ransomware attacks. You can take steps to help protect your storage environment against ransomware and help ensure that you can recover your servers and clients if an attack occurs.
- Securing communications
Your data and passwords are more secure when they are protected by using Secure Sockets Layer (SSL) or Transport Layer Security (TLS), a form of SSL.
- Authenticating IBM Spectrum Protect users by using an LDAP server
Within an IBM Spectrum Protect system, users must authenticate to the server by providing a user ID and password. If your

organization uses a Lightweight Directory Access Protocol (LDAP) server to manage user IDs, you can use the LDAP server to authenticate IBM Spectrum Protect user IDs.

Security concepts

You can protect IBM Spectrum Protect™ from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.

Transport Layer Security

You can use the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol to provide transport layer security for a secure connection between servers, clients, and storage agents. If you send data between the server, client, and storage agent, use SSL or TLS to encrypt the data.

Tip: Any IBM Spectrum Protect documentation that indicates "SSL" or to "select SSL" applies to TLS.

SSL is provided by the Global Security Kit (GSKit) that is installed with the IBM Spectrum Protect server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with a DB2® database instance that is used by any IBM Spectrum Protect servers.

Each server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a certificate authority (CA). You can use your own certificates or purchase certificates from a CA. Either certificate must be installed and added to the key database on the IBM Spectrum Protect server, client, or storage agent. The certificate is verified by the SSL client or server that requests or initiates the SSL communication. Some CA certificates are preinstalled in the key databases, by default.

SSL is set up independently on the IBM Spectrum Protect server, client, and storage agent.

Authority levels

With each IBM Spectrum Protect server, different administrative authority levels are available that determine the tasks that an administrator can complete.

After registration, an administrator must be granted authority by being assigned one or more administrative authority levels. An administrator with system authority can complete any task with the server and assign authority levels to other administrators by using the GRANT AUTHORITY command. Administrators with policy, storage, or operator authority can complete subsets of tasks.

An administrator can register other administrator IDs, grant levels of authority to them, rename IDs, remove IDs, and lock and unlock them from the server.

An administrator can control access to specific client nodes for root user IDs and non-root user IDs. By default, a non-root user ID cannot back up data on the node. Use the UPDATE NODE command to change the node settings to enable backup.

Passwords

By default, the server automatically uses password authentication. With password authentication, all users must enter a password when they access the server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see Authenticating users by using an LDAP server.

Table 1. Password authentication characteristics

Characteristic	More information
Case-sensitivity	Not case-sensitive.
Default password expiration	90 days. The expiration period begins when an administrator ID or client node is first registered to the server. If the password is not changed within this period, the password must be changed the next time that the user accesses the server.

Characteristic	More information
Invalid password attempts	You can set a limit on consecutive invalid password attempts for all client nodes. When the limit is exceeded, the server locks the node.
Default password length	8 characters. The administrator can specify a minimum length. Beginning with Version 8.1.4, the default minimum length for server passwords changed from 0 to 8 characters.

Session security

Session security is the level of security that is used for communication among IBM Spectrum Protect client nodes, administrative clients, and servers and is set by using the SESSIONSECURITY parameter.

The SESSIONSECURITY parameter can be set to one of the following values:

- The STRICT value enforces the highest level of security for communication between IBM Spectrum Protect servers, nodes, and administrators.
- The TRANSITIONAL value specifies that the existing communication protocol is used while you update your IBM Spectrum Protect software to V8.1.2 or later. This is the default. When SESSIONSECURITY=TRANSITIONAL, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

Note: You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Spectrum Protect servers, but you are not required to upgrade servers first. Communication between servers and clients is not interrupted.

For more information about the SESSIONSECURITY parameter values, see the following commands.

Table 2. Commands used to set the SESSIONSECURITY parameter

Entity	Command
Client nodes	<ul style="list-style-type: none"> • REGISTER NODE • UPDATE NODE
Administrators	<ul style="list-style-type: none"> • REGISTER ADMIN • UPDATE ADMIN
Servers	<ul style="list-style-type: none"> • DEFINE SERVER • UPDATE SERVER

Administrators that authenticate by using the DSMADMC command, DSMC command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

Tips:

- Ensure that all IBM Spectrum Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.
- After an administrator authenticates to a V8.1.2 or later server by using a V8.1.2 or later client, the administrator can authenticate only on clients or servers that are using V8.1.2 or later. An administrator command can be issued from any system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Enforce the highest level of security for communication with the IBM Spectrum Protect server by ensuring that all nodes, administrators, and servers use STRICT session security. You can use the SELECT command to determine which servers, nodes, and administrators are using TRANSITIONAL session security and should be updated to use STRICT session security.

Related reference:

Securing communications

SELECT (Perform an SQL query of the database)

Managing administrators

An administrator who has system authority can complete any task with the IBM Spectrum Protect™ server, including assigning authority levels to other administrators. To complete some tasks, you must be granted authority by being assigned one or more authority levels.

Procedure

Complete the following tasks to modify administrator settings.

Task	Procedure
Add an administrator.	<p>To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:</p> <ol style="list-style-type: none"> Register the administrator and specify Pa\$#\$twO as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$twO</pre> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>
Change administrative authority.	<p>Change the authority level for an administrator, ADMIN1.</p> <ul style="list-style-type: none"> Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre> Revoke system authority for the administrator by issuing the following command: <pre>revoke authority admin1 classes=system</pre>
Remove administrators.	<p>Remove an administrator, ADMIN1, from accessing the IBM Spectrum Protect server by issuing the following command:</p> <pre>remove admin admin1</pre>
Temporarily prevent access to the server.	<p>Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.</p>

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Spectrum Protect™.

About this task

By enforcing password authentication and managing password restrictions, you protect your data and your servers from potential security risks.

Procedure

Complete the following tasks to change password requirements for IBM Spectrum Protect servers.

Table 1. Authentication tasks for IBM Spectrum Protect servers

Task	Procedure
Set a limit for invalid password attempts.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details, and then click the Properties tab. Set the number of invalid attempts in the Invalid sign-on attempt limit field. <p>The default value at installation is 0.</p>
Set a minimum length for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of characters in the Minimum password length field.
Set the expiration period for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of days in the Password common expiration field.
Disable password authentication.	<p>By default, the server automatically uses password authentication. With password authentication, all users must enter a password to access the server.</p> <p>You can disable password authentication only for passwords that authenticate with the server (LOCAL). By disabling password authentication, you increase the security risk for the server.</p>
Set a default authentication method.	<p>Issue the SET DEFAULTAUTHENTICATION command. For example, to use the server as the default authentication method, issue the following command:</p> <pre>set defaultauthentication local</pre> <p>To update one client node to authenticate with the server, include AUTHENTICATION=LOCAL in the UPDATE NODE command:</p> <pre>update node authentication=local</pre>

Securing IBM Spectrum Protect on the system

Protect the system where the IBM Spectrum Protect™ server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the directories for the server database and the server instance. Keep the access settings for these directories that you configured during implementation.

- Restricting user access to the server
Authority levels determine what an administrator can do with the IBM Spectrum Protect server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.
- Limiting access through port restrictions
Limit access to the server by applying port restrictions.

Restricting user access to the server

Authority levels determine what an administrator can do with the IBM Spectrum Protect™ server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.

Procedure

1. After you register an administrator by using the REGISTER ADMIN command, use the GRANT AUTHORITY command to set the administrator's authority level. For details about setting and changing authority, see Managing administrators.
2. To control the authority of an administrator to complete some tasks, use the following two server options:
 - a. You can select the authority level that an administrator must have to issue QUERY and SELECT commands with the QUERYAUTH server option. By default, no authority level is required. You can change the requirement to one of the authority levels, including system.
 - b. You can specify that system authority is required for commands that cause the server to write to an external file with the REQSYSAUTHOUTFILE server option. By default, system authority is required for such commands.
3. You can restrict data backup on a client node to only root user IDs or authorized users. For example, to limit backups to the root user ID, issue the REGISTER NODE or UPDATE NODE command and specify the BACKUPINITIATION=root parameter:

```
update node backupinitiation=root
```

Limiting access through port restrictions

Limit access to the server by applying port restrictions.

About this task

You might have to restrict access to specific servers, based on your security requirements. The IBM Spectrum Protect™ server can be configured to listen on four TCP/IP ports: two that can be used for either regular TCP/IP protocols or Secure Sockets Layer (SSL)/Transport Layer Security (TLS) protocols and two that can be used only for the SSL/TLS protocol.

Procedure

You can set the server options to specify the port that you require, as listed in Table 1.

Table 1. Server options and port access

Server option	Port access
TCPPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default value is 1500.
TCPADMINPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default is the value of TCPPORT. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPSPORT options.
SSLTCPSPORT	Specifies the SSL TCP/IP port address for a server. This port listens for SSL-enabled sessions only. A default port value is not available.
SSLTCPADMINPORT	Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. A default port value is not available. Use this option to separate administrative client traffic from regular client traffic that uses the TCPSPORT and SSLTCPSPORT options.

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (SSLTCPSPORT and SSLTCPADMINPORT):

- When you specify the server's SSL-only port for the LLADDRESS on the DEFINE SERVER or UPDATE SERVER command, you must also specify the SSL=YES parameter.
- When you specify the server's SSL-only port for the client's TCPSPORT option, you must also specify YES for the SSL client option.

Protecting the storage environment against ransomware

Storage environments that are connected to the internet can be the target of ransomware attacks. You can take steps to help protect your storage environment against ransomware and help ensure that you can recover your servers and clients if an attack occurs.

About this task

Ransomware is malicious software that is used to gain access to a computer system and encrypt the data. Typically, the initiator of the ransomware attack encrypts data and then contacts the owner of the data to demand a ransom. If the ransom is not paid, the initiator of the attack threatens to leave the data encrypted. For this reason, you can help to protect your storage environment against a ransomware attack by storing a copy of the data at a location that is *not* accessible from the internet.

One possibility is to back up your database to tape and back up clients to copy storage pools on tape, and then transport the tape volumes to a secure, offsite location. If you use this strategy, you can enable the IBM Spectrum Protect™ disaster recovery manager (DRM) function to track the movement of offsite media and register that information in the IBM Spectrum Protect database. DRM consolidates plans, scripts, and other information in a plan file. You can use the plan file to recover your servers and clients after a ransomware attack.

Procedure

1. When you plan your storage environment, consider whether to use tape as a storage medium, and whether to transport the tape volumes offsite. For instructions about setting up tape storage, see [Tape solution](#).
2. When you plan your storage environment, consider whether to use the DRM function to help recover from a ransomware attack, unplanned outage, or disaster. For an introduction to DRM, see [Preparing for and recovering from a disaster by using DRM](#).
3. Review the policies that are set for your storage environment to ensure that enough backup copies are retained, and that the copies are retained for a sufficient number of days. If your newest files are encrypted by ransomware, you can still access previous versions. To set policies, use the Operations Center or the `DEFINE COPYGROUP` and `UPDATE COPYGROUP` commands. For information about preferred settings, see [Retention and expiration of backup versions](#).
4. Monitor your system daily to detect ransomware as soon as possible. For more information, see [Daily monitoring checklist](#) and [Periodic monitoring checklist](#).

Securing communications

Your data and passwords are more secure when they are protected by using Secure Sockets Layer (SSL) or Transport Layer Security (TLS), a form of SSL.

SSL and TLS are the standard technology for creating encrypted sessions between servers and clients. SSL and TLS provide a secure channel for servers and clients to communicate over open communication paths. With SSL and TLS, the identity of the server is verified by using digital certificates.

To protect your storage environment from security threats, servers, clients, and storage agents that use IBM Spectrum Protect™ V8.1.4 or later software are automatically configured to communicate with each other by using the Transport Layer Security (TLS) 1.2 protocol, and self-signed certificates are distributed automatically.

Restrictions pertaining to earlier releases:

- Beginning with IBM Spectrum Protect V8.1.2, SSL is enabled by default for communications between V8.1.2 and later servers and clients. You must manually configure V8.1.2 storage agents to use SSL.
- Storage agents that use V7.1.8 or later software or V8.1.3 or later software are automatically configured to use SSL.

Library clients and library manager servers automatically use SSL to communicate with storage agents that use V8.1.2 or later software or V7.1.8 or later software, but you must manually configure the certificates between them. A storage agent automatically exchanges certificates with its database server.

- Beginning with IBM Spectrum Protect™ Version 8.1.4, you no longer have to manually configure certificates between storage agents, library clients, and library manager servers. Certificates are automatically configured.
- Servers, storage agents, and clients that use IBM Spectrum Protect software versions earlier than V8.1.2 or Tivoli® Storage Manager software versions earlier than V7.1.8 must always be manually configured to use SSL, even if the server or storage agent is using V8.1.3 or later software.

TLS is used for all communication between the server, storage agent, and clients, except when sending or receiving object data. By default, object data is sent and received by using TCP/IP. To improve system performance, use TLS for authentication without encrypting object data. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP

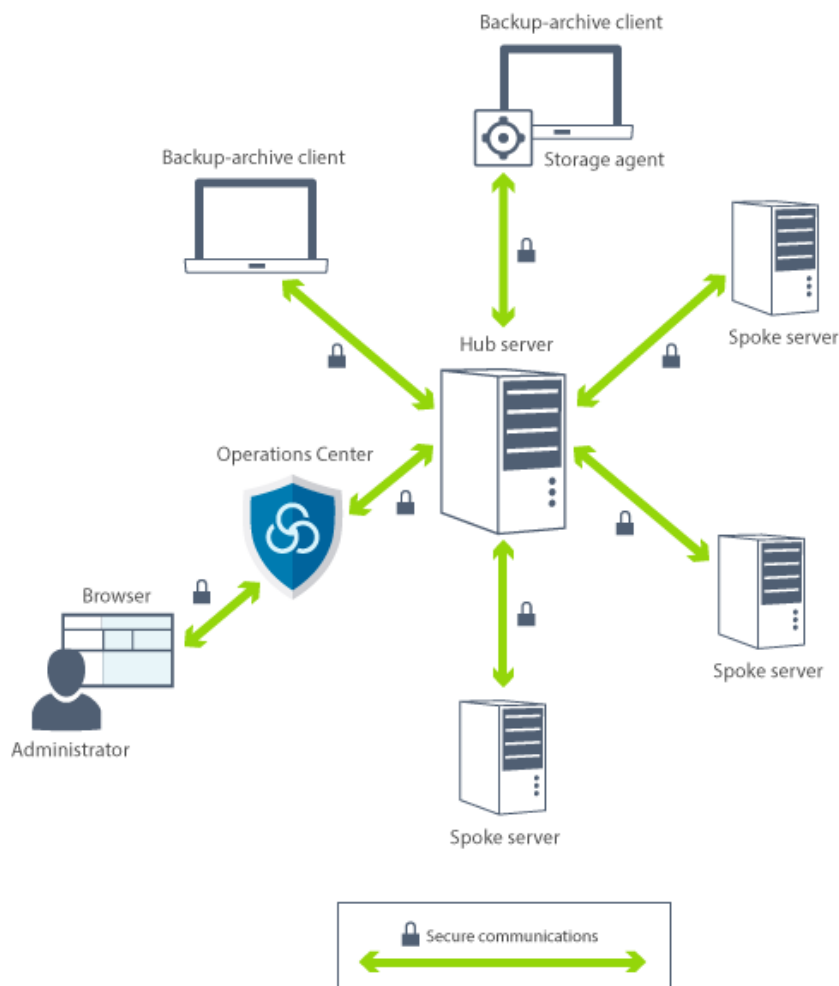
session and the session is secure. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option for client-to-server communication, and the SSL parameter in the UPDATE SERVER command for server-to-server communication. If you choose to use TLS to encrypt object data, consider adding more processor resources on the IBM Spectrum Protect server to manage the increased CPU load.

If you authenticate passwords with an LDAP directory server, TLS protects passwords between the IBM Spectrum Protect server and the LDAP server. TLS is required for all LDAP password communications. Certificates for LDAP directory servers must be manually configured and added to the server key databases. You do not need to add the certificates to storage agent key databases.

- Secure Sockets Layer and Transport Layer Security communication
The Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol is used to provide transport layer security for a secure connection between IBM Spectrum Protect servers, clients, storage agents, and the Operations Center. If you send data between the server, client, and storage agent, SSL or TLS is used to encrypt the data.
- Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL
Configure Secure Sockets Layer (SSL) on the IBM Spectrum Protect server, backup-archive client, storage agent, and the Operations Center to ensure that data is encrypted during communication.

Secure Sockets Layer and Transport Layer Security communication

The Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol is used to provide transport layer security for a secure connection between IBM Spectrum Protect™ servers, clients, storage agents, and the Operations Center. If you send data between the server, client, and storage agent, SSL or TLS is used to encrypt the data.



Restriction: Do not use the SSL or TLS protocols for communications with an IBM DB2® database instance that is used by the IBM Spectrum Protect server.

Each server or storage agent has a unique private key and a unique signed certificate that is used to allow SSL connections. If you use self-signed certificates, the self-signed certificate for each server or storage agent is automatically distributed to all the

clients, storage agents, and servers that use TLS to communicate with it. If you use certificates that are signed by a certificate authority (CA), each IBM Spectrum Protect server and storage agent must send a unique server certificate to a CA to be signed. The CA returns a signed server certificate, which must be added to the server key database, along with the root CA certificate and any intermediate CA certificates. The CA-signed server certificate does not need to be distributed to clients.

If you use CA-signed certificates, all clients, storage agents, and servers that use TLS to communicate with the server or storage agent must have the CA root and intermediate certificates installed in their key database. The CA root and intermediate certificates are used to verify the CA-signed server certificate. The certificates are verified by the SSL client or server that requests or initiates the SSL communication.

The IBM Spectrum Protect server accepts CA-signed certificates that use the SHA-256 or earlier Secure Hash Algorithm encryption method. SHA-256 certificates are designed to improve security and comply with National Institute of Standards and Technology (NIST) requirements. For this reason, the preferred method is to use SHA-256 certificates for communications between the server and Operations Center.

If a server has an MD5-signed certificate that is labeled "TSM Server SelfSigned Key" set as the default when you upgrade to V8.1.4 or later, the default certificate is automatically updated to use a certificate with a SHA signature. In releases prior to V7.1.8, the default certificate was labeled "TSM Server SelfSigned Key" and had an MD5 signature, which does not support the TLS 1.2 protocol that is required by default for V8.1.2 or later clients and the Operations Center. Beginning with V8.1.4, servers that use the MD5-signed certificate as the default are automatically updated to use a default certificate with a SHA signature that is labeled "TSM Server SelfSigned SHA Key". A copy of the certificate is stored in the cert256.arm file, which is located in the server instance directory.

An IBM Spectrum Protect server, client, or storage agent can serve as an SSL client during communication. An SSL client is the component that initiates communication and verifies the certificate for an SSL server. For example, if the IBM Spectrum Protect client initiates the SSL communication with the IBM Spectrum Protect server, the IBM Spectrum Protect client is the SSL client and the server is the SSL server.

Table 1 lists the components that can be an SSL client or SSL server.

Table 1. SSL clients and servers in the IBM Spectrum Protect environment

SSL client	SSL server	Scenario
Client	Server	The IBM Spectrum Protect client initiates a communication request with the IBM Spectrum Protect server. The client verifies the certificate. The server provides the certificate.
Server (such as a source server)	Server (such as a target server)	The IBM Spectrum Protect source server initiates a communication request with the IBM Spectrum Protect target server. The source server acts as an SSL client and verifies the certificate that the target server provides. This type of communication is common during replication processing.
Client through a storage agent	Server	The client verifies each certificate when it initiates SSL communication separately with the IBM Spectrum Protect server and the storage agent. When the storage agent communicates with the server by using the SSL communication protocol, the storage agent acts as an SSL client and verifies the certificate that the server provides. The storage agent can be the SSL client and the SSL provider at the same time. The client must use the same communication protocol (either SSL or TCP/IP) to communicate with both the server and the storage agent.
Server	LDAP server	The IBM Spectrum Protect server initiates a communication request with the LDAP server. The IBM Spectrum Protect server acts as the SSL client and verifies the certificate that the LDAP server provides.
Operations Center	Server	The Operations Center initiates a communication request with the IBM Spectrum Protect server. The Operations Center acts as the SSL client and verifies the certificate that the IBM Spectrum Protect server provides.
Reporting	Server	The reporting agent initiates a communication request with the IBM Spectrum Protect server. The Reporting feature acts as the SSL client and verifies the certificate that the IBM Spectrum Protect server provides.

Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL

Configure Secure Sockets Layer (SSL) on the IBM Spectrum Protect™ server, backup-archive client, storage agent, and the Operations Center to ensure that data is encrypted during communication.

You can use a self-signed SSL certificate or a signed certificate from a third-party certificate authority (CA) to verify an SSL communication request between the server, client, and storage agent. Each IBM Spectrum Protect server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a CA.

The benefit of CA-signed certificates is that a single CA-signed certificate can be used for all servers, which allows you to distribute a single certificate to clients. If you use a self-signed certificate, the certificate is automatically created for each server and storage agent. If you use a root certificate from a CA, it must be installed on each key database for the client, server, and storage agent that initiates SSL communication. The certificate is verified by the SSL client or server that requests or initiates the SSL communication.

Restriction: Some CAs use certificates in a format that is not recognized by IBM Spectrum Protect. You might have to contact your CA to convert the certificate to a format that you can use with IBM Spectrum Protect.

- Configuring the server to accept SSL connections
Configure the server to accept SSL connections before you enable SSL communication from the server to a client, a storage agent, or another server.
- Configuring a storage agent to use SSL
To ensure that data is encrypted for communication between the storage agent and the server and the storage agent and the client, configure the storage agents to communicate by using the SSL protocol.
- Configuring the client to connect to a storage agent by using SSL
To protect the data that is transmitted between a client and storage agent, configure the client to connect to the storage agent by using the SSL protocol.

Configuring the server to accept SSL connections

Configure the server to accept SSL connections before you enable SSL communication from the server to a client, a storage agent, or another server.

About this task

Use this procedure for manual configuration.

Procedure

1. Specify the port on which the server waits for client communications that are enabled for SSL or accept the default port number. Optionally, update the `dsmserv.opt` file in the server instance directory by specifying the `TCP` or `TCPADMIN` options, or both. The `SSLTCP` and `SSLTCPADMIN` options can be used for SSL-only connections.
2. Create the server key database by starting the server. The server key database file, `cert.kdb`, is stored in the server instance directory, and the default certificate label is automatically set as "TSM Server SelfSigned SHA Key". The certificate is exported to the `cert256.arm` file.
3. If you are using the default self-signed certificate, the default self-signed certificate (`cert256.arm`) file is needed when you connect to the server by using TLS. After you use the `cert256.arm` file to import the self-signed certificate to the key database, the file is no longer needed.
4. If you are using a CA-signed certificate, each IBM Spectrum Protect server must send a unique server certificate to a CA to be signed. The CA returns a signed server certificate. To configure CA certificates, complete the following steps for each IBM Spectrum Protect™ server:
 - a. Import the root CA certificate for each IBM Spectrum Protect server that enables SSL. Log on to the IBM Spectrum Protect server system with the instance user ID and issue the following example command from the instance directory:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -label "CA cert" -file ca.crt
```
 - b. Import one or more intermediate CA certificates by issuing the following example command for each intermediate certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -label "Intermediate CA cert" -file
intca.crt
```

- c. The CA root and intermediate certificates (ca.crt and intca.crt) are used to verify the CA-signed server certificate. The CA root and intermediate certificates must be installed in the key database of all clients, storage agents, and servers that use TLS to communicate with the server.
- d. On the server, create a certificate request for the CA to sign by issuing a command that is similar to the following example:

```
gsk8capicmd_64 -certreq -create -db cert.kdb -stashed -label "CA cert"
-sigalg sha256 -size 2048 -ku "digitalSignature,keyEncipherment,keyAgreement"
-eku "clientAuth,serverAuth" -dn "CN=tucson.example.com,OU=Spectrum Protect,O=IBM"
-san_dnsname tucson.example.com -san_ipaddr 9.11.0.0 -file cert_request.csr
```

- e. To receive the signed certificate and make it the default for communicating with clients, issue the following example command:

```
gsk8capicmd_64 -cert -receive -db cert.kdb -stashed -file cert_signed.crt
-default_cert yes
```

The CA-signed server certificate does not need to be distributed to clients.

5. If you made any changes, restart the server.

What to do next

Enable SSL communication from a client, a storage agent, or another server to this server. To complete the following tasks, you must have the server's certificate and the port number that is defined for the server.

1. To enable SSL communication from a client to this server, see [Configuring IBM Spectrum Protect client/server communication with Secure Sockets Layer](#).
 2. To enable SSL communication from another server to this server, see [Configuring the server to connect to another server by using SSL](#).
 3. To enable SSL communication from a storage agent to this server, see [Configuring a storage agent to use SSL](#).
 4. To enable SSL communication from the Operations Center to this server, see [Configuring the Operations Center to connect to the hub server by using SSL](#).
 5. To enable SSL communication from the Data Protection for VMware vSphere GUI to this server, see [Configuring the Data Protection for VMware vSphere GUI to communicate with the server by using SSL](#).
- [Configuring clients to communicate with the server by using SSL](#)
To ensure that data is encrypted during client/server communication, configure clients to communicate with the server by using the SSL protocol.
 - [Configuring the server to connect to another server by using SSL](#)
To ensure that data is encrypted for server-to-server communication, configure servers to communicate with servers by using the SSL protocol.
 - [Configuring the Operations Center to connect to the hub server by using SSL](#)
To ensure that data is encrypted for communication between the Operations Center and the hub server, configure the Operations Center to communicate with the hub server by using the SSL protocol.
 - [Configuring the Data Protection for VMware vSphere GUI to communicate with the server by using SSL](#)
To ensure that data is encrypted while communicating with the IBM Spectrum Protect server, configure the Data Protection for VMware vSphere GUI to communicate with the server by using the SSL protocol.

Related reference:

TCPPORT
TCPADMINPORT
QUERY SESSION (Query client sessions)

Configuring clients to communicate with the server by using SSL

To ensure that data is encrypted during client/server communication, configure clients to communicate with the server by using the SSL protocol.

Before you begin

You must have the server's certificate and the port number that the server is using. For more information, see [Configuring the server to accept SSL connections](#).

Procedure

To enable SSL communication between the server and clients, see [Configuring IBM Spectrum Protect client/server communication with Secure Sockets Layer](#).

Configuring the server to connect to another server by using SSL

To ensure that data is encrypted for server-to-server communication, configure servers to communicate with servers by using the SSL protocol.

Before you begin

You must have the certificate and the port number for the server that you are connecting to. For more information, see [Configuring the server to accept SSL connections](#).

About this task

Tips:

- If both servers are using IBM Spectrum Protect™ V8.1.2 or later software, SSL is automatically configured. Manual configuration is recommended but not required. If either server is using IBM Spectrum Protect software earlier than V8.1.2 or Tivoli® Storage Manager software earlier than V7.1.8, you must manually configure SSL.
- In V8.1.2, you must manually configure storage agents to use SSL. In V8.1.3, storage agents are automatically configured to use SSL.

In the procedure, the following server addresses are used as examples:

- ServerA (the server you are connecting to) is at `bfa.tucson.example.com`
- ServerB is at `bfب.tucson.example.com`

Procedure

1. Create the server key database by starting the server. The server key database file, `cert.kdb`, is stored in the server instance directory.
2. For each server, import the other server's self-signed certificate file (`cert256.arm`) or CA-certificate files. To import the self-signed certificate, issue the following command:

```
gsk8capicmd_64 -cert -add -label server_ip_address -db cert.kdb -stashed  
-file cert256.arm
```

Tip: Use the IP address of the server as the label name.

3. From each server, you can view the certificates in the key database by issuing the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

4. Restart the servers.
5. Issue the DEFINE SERVER command.
 - a. For ServerA, issue the following command:

```
DEFINE SERVER BFB hla=bfب.tucson.example.com lla=1542  
serverpa=passwordforbfب SSL=YES
```

- b. For ServerB, issue the following command:

```
DEFINE SERVER BFA hla=bfa.tucson.example.com lla=1542  
serverpa=passwordforbfa SSL=YES
```

Related reference:

QUERY SESSION (Query client sessions)
TCPPOINT
TCPADMINPORT
DEFINE SERVER (Define a server for server-to-server communications)

Configuring the Operations Center to connect to the hub server by using SSL

To ensure that data is encrypted for communication between the Operations Center and the hub server, configure the Operations Center to communicate with the hub server by using the SSL protocol.

Before you begin

You must have the hub server's certificate and the port number that the server is using. For more information, see [Configuring the server to accept SSL connections](#).

Procedure

To configure SSL communications with the Operations Center, see [Securing communications between the Operations Center and the hub server](#).

Configuring the Data Protection for VMware vSphere GUI to communicate with the server by using SSL

To ensure that data is encrypted while communicating with the IBM Spectrum Protect server, configure the Data Protection for VMware vSphere GUI to communicate with the server by using the SSL protocol.

Before you begin

You must have the server's certificate and the port number that the server is using. For more information, see [Configuring the server to accept SSL connections](#).

Procedure

To enable SSL communication between the server and the Data Protection for VMware vSphere GUI, see [Enabling secure communication with the IBM Spectrum Protect server](#).

Configuring a storage agent to use SSL

To ensure that data is encrypted for communication between the storage agent and the server and the storage agent and the client, configure the storage agents to communicate by using the SSL protocol.

Before you begin

You must have the server's certificate and the port number that the server is using. For more information, see [Configuring the server to accept SSL connections](#).

Procedure

1. Initialize the storage agent and add communication information to the device configuration file and the storage agent options file `dsmsta.opt` by issuing the `DSMSTA SETSTORAGESERVER` command. You must specify the `SSL=YES` parameter to create the key database file in `dsmsta.opt`. All passwords are encrypted in `dsmsta.opt`.

```
dsmsta setstorageserver myname=storage_agent_name mypa=sta_password  
myhla=ip_address servername=server_name serverpa=server_password hla=ip_address lla=ssl_port  
ssl=yes
```

2. Create the key database certificate and default certificates by starting the storage agent.
3. For the storage agent and the server, import the other's `cert256.arm` or CA-certificate files:

```
gsk8capicmd_64 -cert -add -label ip_address -db cert.kdb -stashed  
-file cert256.arm
```

Tip: Use the IP address as the label name.

4. You can view the certificates in the key database by issuing the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

5. Restart the storage agent and the server.
6. Establish communication between the server and the storage agent by issuing the following command:

```
define server sta hla=ip_address lla=port serverpa=password ssl=yes
```

Related reference:

QUERY SESSION (Query client sessions)

TCPPORT

TCPADMINPORT

DEFINE SERVER (Define a server for server-to-server communications)

Configuring the client to connect to a storage agent by using SSL

To protect the data that is transmitted between a client and storage agent, configure the client to connect to the storage agent by using the SSL protocol.

Before you begin

You must have the certificate and the port number for the storage agent.

About this task

After you configure a storage agent to accept SSL connections, configure clients to connect to the storage agent by using SSL.

Procedure

To enable SSL communication between the clients and the storage agent, see [Configuring IBM Spectrum Protect client/server communication with Secure Sockets Layer](#).

Related reference:

TCPPORT

TCPADMINPORT

Authenticating IBM Spectrum Protect users by using an LDAP server

Within an IBM Spectrum Protect™ system, users must authenticate to the server by providing a user ID and password. If your organization uses a Lightweight Directory Access Protocol (LDAP) server to manage user IDs, you can use the LDAP server to authenticate IBM Spectrum Protect user IDs.

You can use one of the following methods to authenticate users with an LDAP server:

Method that is preferred for IBM® Tivoli® Storage Manager Version 7.1.7 and later, and for IBM Spectrum Protect V8.1 and later servers

To use this method, sometimes known as *integrated mode*, user IDs must be registered in an Active Directory database on an LDAP server. Then, you register the same users with the IBM Spectrum Protect server. When a registered user ID accesses the IBM Spectrum Protect server, the credentials are authenticated against the Active Directory database.

To use this method, follow the instructions in [Authenticating users by using an Active Directory database](#).

Method that is used for servers earlier than V7.1.7, and by IBM Security Directory Server users

To use this method, user IDs must be registered in an Active Directory database on an LDAP server. Alternatively, user IDs can be registered in an IBM Security Directory Server (formerly IBM Tivoli Directory Server) database on an LDAP server. With this method, you cannot use the standard user accounts that are registered with the LDAP server. You must create separate user accounts that are associated with a specific organizational unit. To use this method, follow the instructions in [Managing passwords and logon procedures \(V7.1.1\)](#).

- **Authenticating users by using an Active Directory database**
You can authenticate IBM Spectrum Protect users by using an Active Directory database on a Lightweight Directory Access Protocol (LDAP) server. With this method, you use the standard user accounts that are registered with the LDAP server. The same user ID can be used to authenticate to the IBM Spectrum Protect server and to the LDAP server.

Replicating client data to another server

Replicating client data from a source server to another server helps to ensure that backed-up client data is available for recovery if the source server is damaged. Replication incrementally copies data from the source server to the target server to provide failover and failback capability.

About this task

If a disaster occurs and the source server is temporarily unavailable, client nodes can recover their data from the target server. If the source server cannot be recovered, you can change client node configurations to store data on the target server. When an outage occurs, the source server can automatically fail over to a target server for data recovery.

Restriction: A server can replicate data to only one target server.

You can replicate data that is stored in any type of storage pool. The storage pool type can be different at the source replication server and the target replication server. You can control replication by type of client node data:

- Active and inactive backup data together, or only active backup data
- Archive data
- Data that was migrated to a source server by IBM Spectrum Protect™ for Space Management clients

When you replicate data in directory-container storage pools, use storage pool protection to improve the efficiency of the replication process, and to enable repair of data. When you use the Operations Center to set up your storage pools, schedules for protection are automatically defined to coordinate with the replication schedule.

Procedure

1. Verify that servers are compatible and have the system resources for successful use of replication.

Increased amounts of memory and processor cores are required. The database and its logs must be sized to ensure that transactions can complete. A dedicated network, with enough bandwidth to handle the amount of data you intend to replicate, is required.

- a. Verify that the source and target servers are compatible for replication. See [Replication compatibility](#).
- b. Verify that the server has appropriate resources for good performance. For details, see [Checklist for node replication](#).

2. Enable replication. See [Enabling node replication](#).
 3. Schedule replication for the source server. For information about how to integrate this schedule in regular server maintenance schedules, see [Defining schedules for server maintenance activities](#).
 4. Schedule storage pool protection for all directory-container storage pools on the source server. See [Protecting data in directory-container storage pools](#).
 5. Monitor replication by using the Operations Center. For more information, see [Daily monitoring checklist](#).
- **Replication compatibility**
Before you set up replication operations with IBM Spectrum Protect, you must ensure that the source and target replication servers are compatible for replication.
 - **Enabling node replication**
You can enable node replication to protect your data.
 - **Protecting data in directory-container storage pools**
Protect data in directory-container storage pools to reduce node replication time and to enable repair of data in directory-container storage pools.
 - **Modifying replication settings**
Modify replication settings in the Operations Center. Change settings such as the number of replication sessions, replication rules, the data that you want to replicate, the replication schedule, and the replication workload.
 - **Setting different retention policies for the source server and target server**
You can set policies on the target replication server that manage the replicated client-node data differently than on the source server. For example, you can maintain a different number of versions of files on the source and the target servers.

Replication compatibility

Before you set up replication operations with IBM Spectrum Protect™, you must ensure that the source and target replication servers are compatible for replication.

Table 1. Replication compatibility of server versions

Source replication server version	Compatible versions for the target replication server
V7.1	V7.1 or later
V7.1.1	V7.1 or later
V7.1.3	V7.1.3 or later
V7.1.4	V7.1.3 or later
V7.1.5	V7.1.3 or later
V7.1.6	V7.1.3 or later
V7.1.7	V7.1.3 or later
V7.1.8	V7.1.3 or later
V8.1	V7.1.3 or later
V8.1.1	V7.1.3 or later
V8.1.2	V7.1.3 or later
V8.1.3	V7.1.3 or later
V8.1.4	V7.1.3 or later
V8.1.5	V7.1.3 or later

Enabling node replication

You can enable node replication to protect your data.

Before you begin

Ensure that the source and target servers are compatible for replication.

About this task

Replicate the client node to replicate all client data, including metadata. By default, node replication is disabled when you start the server for the first time.

Tips:

- To reduce replication processing time, protect the storage pool before you replicate client nodes. When node replication is started, the data extents that are already replicated through storage pool protection are skipped.
- Replication requires increased amounts of memory and sufficient bandwidth to complete processing. Size the database and its logs to ensure that transactions can complete.

Procedure

To enable node replication, complete the following steps in the Operations Center:

- On the Servers page, click Details.
- On the Details page, click Properties.
- In the Replication section, select Enabled in the Outbound replication field.
- Click Save.

What to do next

Complete the following actions:

1. To verify that replication was successful, review the Daily monitoring checklist.
2. **Linux** If the IBM Spectrum Protect™ server replicates nodes to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Follow the instructions in Determining whether Aspera FASP technology can optimize data transfer in your system environment.

Protecting data in directory-container storage pools

Protect data in directory-container storage pools to reduce node replication time and to enable repair of data in directory-container storage pools.

Before you begin

Ensure that at least one directory-container storage pool exists on the target replication server. When you enable replication in the Operations Center, you can schedule storage pool protection. To configure replication and enable storage pool protection, complete the following steps:

1. On the Operations Center menu bar, hover over Storage and click Replication.
2. On the Replication page, click Server Pair.
3. Complete the steps in the Add Server Pair wizard.

About this task

Protecting a directory-container storage pool backs up data extents to another storage pool, and can improve performance for node replication. When node replication is started, the data extents that are already backed up through storage pool protection are skipped, which reduces the replication processing time. You can schedule the protection of storage pools several times a day to keep up with changes to data.

By protecting a storage pool, you do not use resources that replicate existing data and metadata, which improves server performance. You must use directory-container storage pools if you want to protect and back up the storage pool only.

Alternative protection strategy: As an alternative to using replication, you can protect data in directory-container storage pools by copying the data to container-copy storage pools. Data in container-copy storage pools is stored on tape volumes. Tape copies that are stored offsite provide additional disaster recovery protection in a replicated environment.

Procedure

1. Alternatively, to enable storage pool protection, you can use the PROTECT STGPOOL command from the source server to back up data extents in a directory-container storage pool. For example, to protect a directory-container storage pool that is named POOL1 issue the following command:

```
protect stgpool pool1
```

As part of the operation of the PROTECT STGPOOL command, damaged extents in the target storage pool are repaired. To be repaired, extents must already be marked as damaged on the target server. For example, an AUDIT CONTAINER command might identify damage in the target storage pool before the PROTECT STGPOOL command is issued.

2. Optional: If damaged extents were repaired in the target storage pool and you protect multiple source storage pools in one target storage pool, complete the following steps to ensure a complete repair:
 - a. Issue the PROTECT STGPOOL command for all source storage pools to repair as much of the damage as possible.
 - b. Issue the PROTECT STGPOOL command again for all source storage pools. For this second operation, use the FORCERECONCILE=YES parameter. This step ensures that any repairs from other source pools are properly recognized for all source storage pools.


Results

If a directory-container storage pool is protected, you can repair the storage pool if damage occurs, by using the REPAIR STGPOOL command.

Restriction: If you replicate client nodes but do not protect the directory-container storage pool, you cannot repair the storage pool.

What to do next

Complete the following actions:

1. To view replication workload status, follow the instructions in the Daily monitoring checklist.
2.  If the IBM Spectrum Protect™ server replicates nodes to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Follow the instructions in Determining whether Aspera FASP technology can optimize data transfer in your system environment.

Related tasks:

[Copying directory-container storage pools to tape](#)

Related reference:

[AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

[PROTECT STGPOOL](#) (Protect storage pool data)

Modifying replication settings

Modify replication settings in the Operations Center. Change settings such as the number of replication sessions, replication rules, the data that you want to replicate, the replication schedule, and the replication workload.

About this task

You might need to customize your replication settings in the following scenarios:

- Changes to data priorities
- Changes to replication rules
- Requirement for a different server to be the target server
- Scheduled processes that negatively affect server performance

Procedure

Use the Operations Center to modify replication settings.

Task	Procedure
Change a replication rule.	<ol style="list-style-type: none"> On the Servers page, click Details. On the Details page, click Properties. In the Replication section, choose the replication rule that you want to apply: Default archive rule, Default backup rule, or Default space-management rule. Click Save.
Specify the duration that replication records are retained.	<ol style="list-style-type: none"> On the Servers page, click Details. On the Details page, click Properties. In the Replication section, enter the number of days that replication records must be retained in the Retain replication history field. Alternatively, select the Do not retain check box if you do not require replication records. Click Save.
Specify a target replication server.	<ol style="list-style-type: none"> On the Servers page, click Details. On the Details page, click Properties. In the Replication section, specify the target server. Click Save.
Cancel a replication process.	<ol style="list-style-type: none"> On the Servers page, click Active tasks. Select the process or session that you want to cancel. Click Cancel.

Setting different retention policies for the source server and target server

You can set policies on the target replication server that manage the replicated client-node data differently than on the source server. For example, you can maintain a different number of versions of files on the source and the target servers.

Procedure

1. From the source replication server, validate the replication configuration and verify that the source replication server can communicate with the target replication server by issuing the `VALIDATE REPLICATION` command. For example, validate the configuration by using the name of one client node that is being replicated:

```
validate replication node1 verifyconnection=yes
```

2. From the source replication server, issue the `VALIDATE REPLPOLICY` command to review the differences between the policies on the source and target replication servers. For example, to display the differences between the policies on the source server and the target server, `CVT_SRV2`, issue the following command from the source server:

```
validate replpolicy cvt_srv2
```

3. Update the policies on the target server if necessary.

Tip: You can use the Operations Center to modify the policies on the target server. Follow the instructions in [Editing policies](#).

For example, to maintain inactive versions of files for a shorter time on the target server than on the source server, reduce the Backups setting in the management classes that apply to replicated client data.

4. Enable the target replication server to use its policies to manage the replicated client-node data by issuing the `SET DISSIMILARPOLICIES` command on the source server. For example, to enable the policies on the target replication server, `CVT_SRV2`, issue the following command on the source server:

```
set dissimilarpolicies cvt_srv2 on
```

The next time that the replication process runs, the policies on the target replication server are used to manage the replicated client-node data.

Tip: If you configure replication by using the Operations Center and the policies on the source and target replication servers do not match, the policy that is specified for the source replication server is used. If you enabled the policies on the target replication server by using the `SET DISSIMILARPOLICIES` command, the policy that is specified for the target replication server is used. If the target replication server does not have the policy that is used by the node on the source replication server, the `STANDARD` policy is used.

Related reference:

[EXPORT POLICY](#) (Export policy information)

[SET DISSIMILARPOLICIES](#) (Enable the policies on the target replication server to manage replicated data)

[VALIDATE REPLICATION](#) (Validate replication for a client node)

[VALIDATE REPLPOLICY](#) (Verify the policies on the target replication server)

AIX

Linux

Windows

Configuring clustered environments

You can configure the IBM Spectrum Protect™ server for clustering on AIX®, Linux, or Windows systems.

You can use a clustered environment for the following operating systems:

- IBM® PowerHA® SystemMirror for AIX
- IBM Tivoli® System Automation for Multiplatforms for AIX and Linux
- Microsoft Failover Cluster for Windows

You can use other cluster products with IBM Spectrum Protect, however, documentation is not available and support is limited. For the latest information about support for clustered environments, see <http://www.ibm.com/support/docview.wss?uid=swg21609772>.

Before you use another cluster product, verify that DB2® supports the required file systems. For more information about the level of DB2 that you are using, see the DB2 product information, and search for recommended file systems.

- [Clustered environment overview](#)
Clusters consist of many components such as IBM Spectrum Protect servers, hardware, and software. You can use clustering to join two or more servers or nodes by using a shared disk system.
- [AIX](#) [Configuring an AIX environment for clustering](#)
You can configure the IBM Spectrum Protect server for AIX clustered environments by using IBM PowerHA SystemMirror for AIX or IBM Tivoli System Automation for Multiplatforms.
- [Linux](#) [Configuring a Linux environment for clustering](#)
You can configure the IBM Spectrum Protect Linux server in a clustered environment by using IBM Tivoli System Automation for Multiplatforms Version 4.1.

- **Windows** Configuring a Windows clustered environment

You can configure an IBM Spectrum Protect server for Windows in a Microsoft failover cluster environment. Windows cluster environments consist of components such as IBM Spectrum Protect servers, hardware, and software. When these components are connected to the same disk system, downtime is minimized.

Related information:

Upgrading the server in a clustered environment

AIX

Linux

Windows

Clustered environment overview

Clusters consist of many components such as IBM Spectrum Protect™ servers, hardware, and software. You can use clustering to join two or more servers or nodes by using a shared disk system.

This configuration provides the nodes with the ability to share data, which allows higher server availability and minimized downtime. For example:

- You can configure, monitor, and control applications and hardware components that are deployed on a cluster.
- You can use an administrative cluster interface and IBM Spectrum Protect to designate cluster arrangements and define a failover pattern. The server is part of the cluster that provides an extra level of security by ensuring that no transactions are missed because a server failed. The failover pattern that you establish prevents future failures.
- You can apply clustering to the node replication process. In this way, server availability is higher than it would be if node replication is used as a process on its own. Server availability is higher because a client is less likely to fail over to another server in a clustered environment. If you replicate data from several source replication servers to one target replication server, there is a high dependency on the target replication server. A clustered environment eases the dependency on the target replication server.

Components in a server cluster are known as *cluster objects*. Cluster objects are associated with a set of properties that have data values that describe the identity and behavior of an object in the cluster. Cluster objects can include the following components:

- Nodes
- Storage
- Services and applications
- Networks

You manage cluster objects by manipulating their properties, typically through a cluster management application.

- Cluster nodes
Nodes in a cluster all have similar characteristics, which allows them to work together.

AIX

Configuring an AIX environment for clustering

You can configure the IBM Spectrum Protect™ server for AIX® clustered environments by using IBM® PowerHA® SystemMirror for AIX or IBM Tivoli® System Automation for Multiplatforms.

PowerHA SystemMirror for AIX and System Automation for Multiplatforms detect system failures and manage failover to a recovery processor with a minimal loss of user time. You can set up the IBM Spectrum Protect server on a system in a PowerHA or a System Automation for Multiplatforms cluster. Then, if the system fails, the IBM Spectrum Protect server can be started on another system in the cluster.

In both failover and failback, it seems that the IBM Spectrum Protect server halted and is then restarted. Any transactions that were in progress at the time of the failover or failback are rolled back, and all completed transactions are still complete. IBM Spectrum Protect clients see failover or failback as a communications failure and try to reestablish their connections.

See the following information for details about these clustering options.

- Configure IBM Spectrum Protect for AIX to use IBM PowerHA SystemMirror for AIX in a clustered environment by reviewing the following topics.
- Configure IBM Spectrum Protect for AIX to use System Automation for Multiplatforms in a clustered environment by reviewing the information in <http://www.ibm.com/support/docview.wss?uid=swg27039780>.
- Learn more about PowerHA SystemMirror® product information.

- AIX Requirements for a PowerHA cluster
 IBM PowerHA SystemMirror for AIX detects system failures and manages failover to a recovery processor with a minimal loss of user time.
- AIX PowerHA failover and failback
 If a node fails, the server cluster transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called *failover*. The reverse process, *failback*, occurs when the failed node becomes active again and the groups that were failed over to the other nodes are transferred back to the original node.
- AIX Installing and configuring PowerHA SystemMirror for AIX
 You can configure the IBM Spectrum Protect server for AIX clustered environments by using IBM PowerHA SystemMirror for AIX.
- AIX Installing the IBM Spectrum Protect server on a production node for PowerHA
 Install the IBM Spectrum Protect server on a production node for PowerHA to be able to configure the server for clustering.
- AIX Installing the IBM Spectrum Protect client on a production node for PowerHA
 You need to install only the backup-archive client file set, which contains the backup-archive client files and the administrative command-line client.
- AIX Verifying the configuration of the IBM Spectrum Protect server for PowerHA
 When you configure the IBM Spectrum Protect server to use PowerHA, you must verify the configuration.
- AIX Setting up the standby node for PowerHA
 For PowerHA, ensure that the IBM Spectrum Protect server is not running on the production node before you set up the standby node.
- AIX Defining the removable media storage devices to AIX for PowerHA
 For an AIX operating system, you must define the removable-media storage devices that are used by IBM Spectrum Protect on the production and standby nodes. The library manager validates that the cartridge that contains the removable-media storage device is in the correct drive.
- AIX Completing the cluster manager and IBM Spectrum Protect configurations
 Update the cluster manager configuration to define the IBM Spectrum Protect server as an application and a failover resource of the standby node. This application is owned by the production node.
- AIX Troubleshooting the PowerHA clustered environment
 Review the following list for information about troubleshooting common problems. The information that is provided for IBM PowerHA SystemMirror for AIX does not represent all possible scenarios.

AIX

Requirements for a PowerHA cluster

IBM PowerHA® SystemMirror for AIX® detects system failures and manages failover to a recovery processor with a minimal loss of user time.

The following hardware requirements are for configuring the IBM Spectrum Protect™ server:

- A hardware configuration that is suitable for PowerHA. The removable media storage devices for the IBM Spectrum Protect server must be physically connected to at least two nodes of the PowerHA cluster on a shared bus (including a SAN).
- Sufficient shared disk space to hold the IBM Spectrum Protect database, recovery logs, instance directory, and disk storage pools to be used. See *Managing inventory capacity* to determine how much space is required for the database and recovery log and to ensure the availability of the database and recovery log.
- A TCP/IP network.

Tip: If an IBM Spectrum Protect server manages removable media storage devices, you can configure two IBM Spectrum Protect servers to run on different systems in an PowerHA cluster. Either system can run both servers if the other system fails. To configure two IBM Spectrum Protect servers to run on different systems in an PowerHA cluster, use another file system that is accessible to both servers.

AIX

PowerHA failover and failback

If a node fails, the server cluster transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called *failover*. The reverse process, *failback*, occurs when the failed node becomes active again and the groups that were failed over to the other nodes are transferred back to the original node.

The terms *production node* and *standby node* refer to the two PowerHA® nodes on which the IBM Spectrum Protect™ server runs.

PowerHA manages taking over the TCP/IP address and mounting the shared file system on the standby node or production node, as appropriate.

When a *failover* or *failback* occurs, any transactions that were being processed at the time are rolled back. To IBM Spectrum Protect clients, *failover* or *failback* represents a communications failure. Therefore, you must reestablish a connection that is based on their COMMRESTARTDURATION and COMMRESTARTINTERVAL option settings.

Typically, you can restart the backup-archive client from the last committed transaction. If a client schedule is running when a *failover* occurs, the client operation is likely to fail. If you can restart client operations, then you must restart them from the beginning of the processing. The clients and agent operations complete as they normally do if the server was halted and restarted while they were connected. The only difference is that the server is physically restarted on different hardware.

If you do not want automatic *failback* to occur, you can configure the resource as a cascading resource group without *failback*.

Related information:

[PowerHA SystemMirror product information](#)

Installing and configuring PowerHA SystemMirror for AIX

You can configure the IBM Spectrum Protect™ server for AIX® clustered environments by using IBM® PowerHA® SystemMirror® for AIX.

- **AIX** Installing and configuring the PowerHA cluster
You might experience processing errors if your IBM PowerHA SystemMirror for AIX installation and configuration are not done correctly.
- **AIX** Configuring IBM Spectrum Protect server on the primary node for PowerHA
You can configure an IBM Spectrum Protect server instance on the primary node.
- **AIX** Configuring IBM Spectrum Protect server on a secondary node for PowerHA with a shared DB2 instance
If the DB2® instance directory is shared between the nodes in the PowerHA cluster, you do not need to create a DB2 instance on the secondary node. You do not run the dsmdirfx wizard.
- **AIX** Configuring IBM Spectrum Protect server on a secondary node for PowerHA with a separate DB2 instance
You must create a DB2 instance on each secondary node if the DB2 instance directory, /home/tsminst1/sqllib, is not shared between the nodes in the PowerHA cluster.

AIX

Installing and configuring the PowerHA cluster

You might experience processing errors if your IBM PowerHA® SystemMirror for AIX® installation and configuration are not done correctly.

Procedure

Complete the following steps to install and configure the PowerHA cluster:

1. Define the shared file systems and logical volumes when they are needed. You might want to put files in separate file systems or on separate physical disks for integrity or performance reasons. Do not put the home directory of the user instance on a shared disk. Mirror the logical volumes to provide maximum availability (including the underlying file systems). The file systems that must be defined include the IBM Spectrum Protect™ server instance directory, the database and log directories, all disk storage pool directories, and FILE device type storage pool directories.
2. Configure PowerHA so that the production node owns the shared volume groups and the standby node takes over the shared volume groups if the production node fails.
3. Configure PowerHA so that the file systems also fail over.
4. Set up a Service IP address for the IBM Spectrum Protect server. The Service IP address must be different from each host IP address. The Service IP is moved from host to host, not the actual host IP address.
5. Failover the shared database and the log and instance directories to the standby node of the PowerHA cluster.

Results

You must configure the removable media storage devices for failover and define the IBM Spectrum Protect server as an application to PowerHA.

AIX

Configuring IBM Spectrum Protect server on the primary node for PowerHA

You can configure an IBM Spectrum Protect™ server instance on the primary node.

Procedure

1. Review the topics in the configuring the IBM Spectrum Protect server information.
2. After you configure the IBM Spectrum Protect server instance on the primary node, you can configure the IBM Spectrum Protect server on a secondary node.

Related tasks:

Configuring the IBM Spectrum Protect server instance

AIX

Configuring IBM Spectrum Protect server on a secondary node for PowerHA with a shared DB2 instance

If the DB2® instance directory is shared between the nodes in the PowerHA® cluster, you do not need to create a DB2 instance on the secondary node. You do not run the dsmicfgx wizard.

Procedure

To configure a server instance on the secondary node with a shared DB2 instance, complete the following steps:

1. On each node in the cluster, add the following text to the `/opt/tivoli/tsm/server/bin/rc.dsmserv` script:

```
DB2NODES_TEMP='/tmp/db2nodes.tmp'
DB2NODES=${homeDir}/sqllib/db2nodes.cfg
# Current hostname
HOSTNAME=$(/bin/hostname)
# hostname saved in db2nodes.cfg
DB2_HOST=$(cat $DB2NODES | cut -d ' ' -f 2)
# if they are different update the file
if [[ "$HOSTNAME" != "$DB2_HOST" ]]
then
  echo "Updating hostname in db2nodes.cfg"
  sed -e s_${DB2_HOST}_${HOSTNAME}_g $DB2NODES > $DB2NODES_TEMP
  cp $DB2NODES_TEMP $DB2NODES
fi
```

Tip: If the text is not included in the script, you can include it before you issue `/opt/tivoli/tsm/server/bin/rc.dsmserv` script.

2. Move all the shared resources to the secondary node.
3. Update the following variables in the `/opt/tivoli/tsm/server/bin/startserver` script, by using the following values:

Table 1. Variables in the `/opt/tivoli/tsm/server/bin/startserver` script

Description	Variable	Example
Set the INST_USER to the instance user ID.	INST_USER	INST_USER='tsmuser1'
Set the INST_DIR to the location of the IBM Spectrum Protect™ instance directory. This directory contains dsm serv.dbid and dsm serv.opt.	INST_DIR	INST_DIR='/home/tsm user1/tsminst1'
Select one of the following startup options: Option 1 - use instance: \$INST_USER but run the server as root (-U) Option 2 - use instance: \$INST_USER and run the server as \$INST_USER (-u)	INST_OPTION	Option 1: INST_OPTION='-U \$INST_USER' Option 2: INST_OPTION='-u \$INST_USER'

4. Start the server by issuing the following script:

```
/opt/tivoli/tsm/server/bin/startserver
```

5. When the server is started, issue the BACKUP DB command to verify that the data is successfully backed up.

AIX

Configuring IBM Spectrum Protect server on a secondary node for PowerHA with a separate DB2 instance

You must create a DB2® instance on each secondary node if the DB2 instance directory, /home/tsminst1/sqllib, is not shared between the nodes in the PowerHA® cluster.

About this task

You can configure the IBM Spectrum Protect™ server on a secondary node by using the dsmsicfgx wizard or manually.

Procedure

- To create a DB2 instance on a secondary node by using the dsmsicfgx wizard, complete the following steps:
 - Run the dsmsicfgx wizard.
 - From the Instance Directory panel, select the Check this if you are configuring the server instance on a secondary node of a high availability cluster check box.
- To create a DB2 instance on a secondary node manually, complete the following steps:
 - Move all the shared resources to the secondary node.
 - Create a DB2 instance by issuing the following db2icrt command:

```
/opt/tivoli/tsm/db2/instance/db2icrt -s ese -u instance_user instance_user
```

where *instance_user* is the same user that owns the DB2 instance on the primary node.

- When the DB2 instance is created, log in as the instance user or by issuing the su command:

```
su - <instance_user>
```

- As the instance user, issue the following commands:

```
db2start
db2 update dbm cfg using DFTDBPATH shared_db_path
db2 catalog db TSMDB1
db2stop
```

where *shared_db_path* is the shared database directory. The shared database directory is typically the server instance directory.

Tip: To determine the *shared_db_path* value, issue the following command on the primary node:

```
db2 get dbm cfg | grep DFTDBPATH
```

- Update the following variables in the /opt/tivoli/tsm/server/bin/startserver script, by using the following values:

Table 1. Variables in the /opt/tivoli/tsm/server/bin/startserver script

Description	Variable	Example
Set the INST_USER to the instance user ID.	INST_USER	INST_USER='tsmuser1'
Set the INST_DIR to the location of the IBM Spectrum Protect instance directory. This directory contains dsmserv.dbid and dsmserv.opt.	INST_DIR	INST_DIR='/home/tsmuser1/tsminst1'

Description	Variable	Example
Select one of the following startup options: Option 1 - use instance: \$INST_USER but run the server as root (-U) Option 2 - use instance: \$INST_USER and run the server as \$INST_USER (-u)	INST_OPTION	Option 1: INST_OPTION='-U \$INST_USER' Option 2: INST_OPTION='-u \$INST_USER'

6. Start the server by issuing the following script:

```
/opt/tivoli/tsm/server/bin/startserver
```

7. When the server is started, issue the BACKUP DB command to verify that the data is successfully backed up.

AIX

Installing the IBM Spectrum Protect server on a production node for PowerHA

Install the IBM Spectrum Protect™ server on a production node for PowerHA® to be able to configure the server for clustering.

Procedure

Complete the following steps to install the IBM Spectrum Protect server on the production node:

1. Install IBM Spectrum Protect. Select one of the following components:
 - o The IBM Spectrum Protect server
 - o The IBM Spectrum Protect device driver, if needed
 - o The IBM Spectrum Protect license

The executable files are typically installed on the internal disks of the production node, not on the shared IBM Spectrum Protect disk space. IBM Spectrum Protect server executable files are installed in the /opt/tivoli/tsm/server/bin directory.

2. Configure IBM Spectrum Protect to use the TCP/IP communication method. For instructions, see the information about configuring a server instance in AIX: Taking the first steps after you install IBM Spectrum Protect.
3. Define a new user ID that owns the IBM Spectrum Protect server instance or use an existing user ID that does not already own an IBM Spectrum Protect instance. While logged in to the instance user ID, complete the following steps:
 - a. Create an instance directory by using the mkdir command on a shared file system that can fail over to the standby system. This disk must be defined to PowerHA.
 - b. Create the database and log directories by using the mkdir command on shared file systems that can fail over to the standby system. These disks must also be defined to PowerHA to fail over.
 - c. Complete the configuration by using the dsomicfgx wizard.

Related tasks:

AIX: Installing the server

Upgrading the server

AIX

Installing the IBM Spectrum Protect client on a production node for PowerHA

You need to install only the backup-archive client file set, which contains the backup-archive client files and the administrative command-line client.

Procedure

For detailed instructions on installing the IBM Spectrum Protect™ client, see Installing the IBM Spectrum Protect backup-archive clients.

Complete the following steps to install the IBM Spectrum Protect client on the production node.

1. Install the IBM Spectrum Protect client executable files in the `/usr/tivoli/tsm/client/ba/bin` directory. These files are typically installed on the internal disks of the production node.
2. For the client to find the server, ensure that the client options file, `dsm.sys`, points to the IBM Spectrum Protect server. The server name in `dsm.sys` is used only on the `-servername` parameter of the `dsmadmcmd` command to specify the server to be contacted.

AIX

Verifying the configuration of the IBM Spectrum Protect server for PowerHA

When you configure the IBM Spectrum Protect™ server to use PowerHA®, you must verify the configuration.

About this task

When you use PowerHA, all database, log, storage, and instance directories must be on shared disks that are configured to fail over by PowerHA.

Procedure

To identify the directories that are on shared disk, complete the following steps:

1. Log on as the instance user.
2. Run the `/opt/tivoli/tsm/server/bin/dsmclustfs` script.
3. Examine the file systems that are reported by the script and verify that they are on shared disks. The following example script shows the type of information that you must review:

```
> su - tsminst1
$ /opt/tivoli/tsm/server/bin/dsmclustfs
SQL1026N The database manager is already active.
```

The following database connection information is displayed when the IBM Spectrum Protect server connects to the DB2® database:

```
DB20000I The START DATABASE MANAGER command completed successfully.
```

```
Database Connection Information
```

```
Database server          = DB2/AIX64 11.1.0
SQL authorization ID     = TSMINST1
Local database alias     = TSMDB1
```

```
File systems for the DB2 database: /TSMdbspace2 /TSMdbspace1
File system for Active Log: /TSMalog
File system for Archive Log: /TSMarchlog
Active log mirror not defined for this database
```

The following mandatory DB2 file systems are in the script:

```
/TSMdb-1 /TSMalog-1 /TSMarchlog-1
```

```
Checking existing TSM disk-based volumes...
TSM Data is stored in the following file systems: /TSMdisk-1 /TSMfile-1
```

AIX

Setting up the standby node for PowerHA

For PowerHA®, ensure that the IBM Spectrum Protect™ server is not running on the production node before you set up the standby node.

Procedure

Complete the following steps to set up the standby node:

1. On the standby node, open the shared volume group and any IBM Spectrum Protect file systems.
2. On the standby node, install the IBM Spectrum Protect product code. For more information, see [Installing the IBM Spectrum Protect server on a production node for PowerHA](#). If the executable files are installed on shared disk space, you might need to install them on the standby node. IBM Spectrum Protect device drivers, SMIT panels, and other files must be installed in AIX® system directories.
3. Open the `dsmicfgx` wizard. Follow the instructions to complete the configuration. Select the check box to indicate that this item is a secondary node in the cluster.
4. Start the server on the standby node. Query the database, recovery log, and storage pool volumes to verify that they are the same as when the server was started on the production node.
5. Install the client on the standby node. If the executable files are installed on shared disk space, you might need to install them on the standby node. IBM Spectrum Protect SMIT panels and other files must be installed in AIX system directories. Use the AIX RCP command with the `-p` option to copy the `dsm.sys` file from the production node to the standby node. If the `dsm.sys` file is changed on one node, it must be copied to the other node.

Results

Tip: If the `dsm.sys` file is changed on one node, you must copy it to the other node.

AIX

Defining the removable media storage devices to AIX for PowerHA

For an AIX® operating system, you must define the removable-media storage devices that are used by IBM Spectrum Protect™ on the production and standby nodes. The library manager validates that the cartridge that contains the removable-media storage device is in the correct drive.

About this task

Prerequisite:

- If you define a library manager server that is not shared with the IBM Spectrum Protect server, ensure that the `RESETDRIVES` parameter for the `DEFINE LIBRARY` command or the `UPDATE LIBRARY` command is specified as `YES`. If you define a library manager server that is shared with the IBM Spectrum Protect server, the `SANDISCOVERY` option must be set to `ON` in the IBM Spectrum Protect server option file `dsmserv.opt`. By default, this option is set to `OFF`.
- You can issue the `PERFORM LIBACTION` command from SCSI and VTL library types. Use this command to define the drives and paths for a library in one step.

If your SAN device mapping is accurate, continue to the [Completing the cluster manager and IBM Spectrum Protect configurations](#) section. If the device names on the primary and secondary systems are not the same, you must use SAN discovery so that the IBM Spectrum Protect server can access the devices.

Related tasks:

[Configuring library sharing \(V7.1.1\)](#)

Related reference:

`DEFINE LIBRARY` (Define a library)

`UPDATE LIBRARY` (Update a library)

`PERFORM LIBACTION` (Define or delete all drives and paths for a library)

`SANDISCOVERY`

Related information:

[IBM Spectrum Protect Supported Devices](#)

AIX

Completing the cluster manager and IBM Spectrum Protect configurations

Update the cluster manager configuration to define the IBM Spectrum Protect™ server as an application and a failover resource of the standby node. This application is owned by the production node.

About this task

You can issue IBM® PowerHA® SystemMirror for AIX® or System Automation for Multiplatforms commands to set up the cluster. Continue with configuring the IBM Spectrum Protect server.

Related information:

[PowerHA SystemMirror product information](#)

[IBM Tivoli System Automation for Multiplatforms Version 3.2.2 product information](#)

AIX

Troubleshooting the PowerHA clustered environment

Review the following list for information about troubleshooting common problems. The information that is provided for IBM® PowerHA® SystemMirror for AIX® does not represent all possible scenarios.

Warning messages that are issued after you run the clverify utility

You can run the PowerHA cluster verification utility, clverify, on one node to verify the cluster configuration and the assignment on the PowerHA resources. If you run the clverify utility after you define the IBM Spectrum Protect™ server as a PowerHA application, warning messages are issued.

Warning messages display because the shell scripts that start and stop the IBM Spectrum Protect servers are in a shared file system. The shell scripts can be run only on one node at a time. Therefore, the shell scripts can be available on only one node at a time. You can ignore the clverify utility warning messages. If a shared file system cannot be mounted, the IBM Spectrum Protect server cannot be started.

IBM Spectrum Protect server fails to start after you issue the startserver script

If you use the startserver shell script and PowerHA fails to start the IBM Spectrum Protect server, start it manually on a terminal without the quiet option. If you want to run the server with the quiet option, issue the dsmserv -q command.

Messages that are associated with the tctl command

If you issue the tctl -f/dev/rmt2 rewind command, the following message might be displayed:

```
/dev/rmt2: A device is already mounted or cannot be unmounted
```

This message means that the I/O device is locked with a SCSI RESERVE by a system other than the one on which the tctl command was run. If you are using persistent reservation, the IBM Spectrum Protect server preempts a drive reservation by default. If the device driver does not use persistent reservation, the server completes a target reset.

ANS4329S Server out of data storage space message

If the ANS4329S Server out of data storage space message is displayed on an IBM Spectrum Protect client, the license for the IBM Spectrum Protect server might be non-compliant. Issue the QUERY LICENSE command to display the compliance information for the license. If the compliance state is valid, use the QUERY ACTLOG command on the server and review the messages that are displayed to identify the problem.

Linux

Configuring a Linux environment for clustering

You can configure the IBM Spectrum Protect™ Linux server in a clustered environment by using IBM® Tivoli® System Automation for Multiplatforms Version 4.1.

- **Linux** Overview of a two-node IBM Spectrum Protect cluster using System Automation for Multiplatforms
Use the System Automation for Multiplatforms cluster for higher server and database availability during a failure. By using the System Automation for Multiplatforms failover function, server components such as the database can automatically recover from a failure.
- **Linux** Setting up an IBM Spectrum Protect cluster with System Automation for Multiplatforms
You must set up the IBM Spectrum Protect cluster to use System Automation for Multiplatforms.
- **Linux** Prerequisites for configuring a Linux clustered environment with System Automation for Multiplatforms
Before you install and configure IBM Spectrum Protect in a clustered environment with System Automation for Multiplatforms, you must check the prerequisites.
- **Linux** Installing and configuring IBM Spectrum Protect components on the primary and secondary nodes
You must install the IBM Spectrum Protect server and database components on the primary and secondary nodes in the cluster. Then, configure the primary node first followed by the secondary node.
- **Linux** Installing System Automation for Multiplatforms on the primary and secondary nodes
After you install and configure IBM Spectrum Protect on the primary and secondary nodes in the cluster, you must install and configure System Automation for Multiplatforms on these nodes. Then, you must activate these nodes for the domain,

configure the resources, and activate the base policy. Finally, you must add the mount points to the IBM Spectrum Protect directories.

- **Linux** Configuring storage resources

Use the System Automation for Multiplatforms user interface or command line to add or delete storage resources and to delete mount points that are no longer required. If you add a storage pool to the cluster, you must add it to the resource group. If you remove a storage pool from the cluster, you must also delete it from the resource group.

- **Linux** Upgrading a server that is configured with System Automation for Multiplatforms

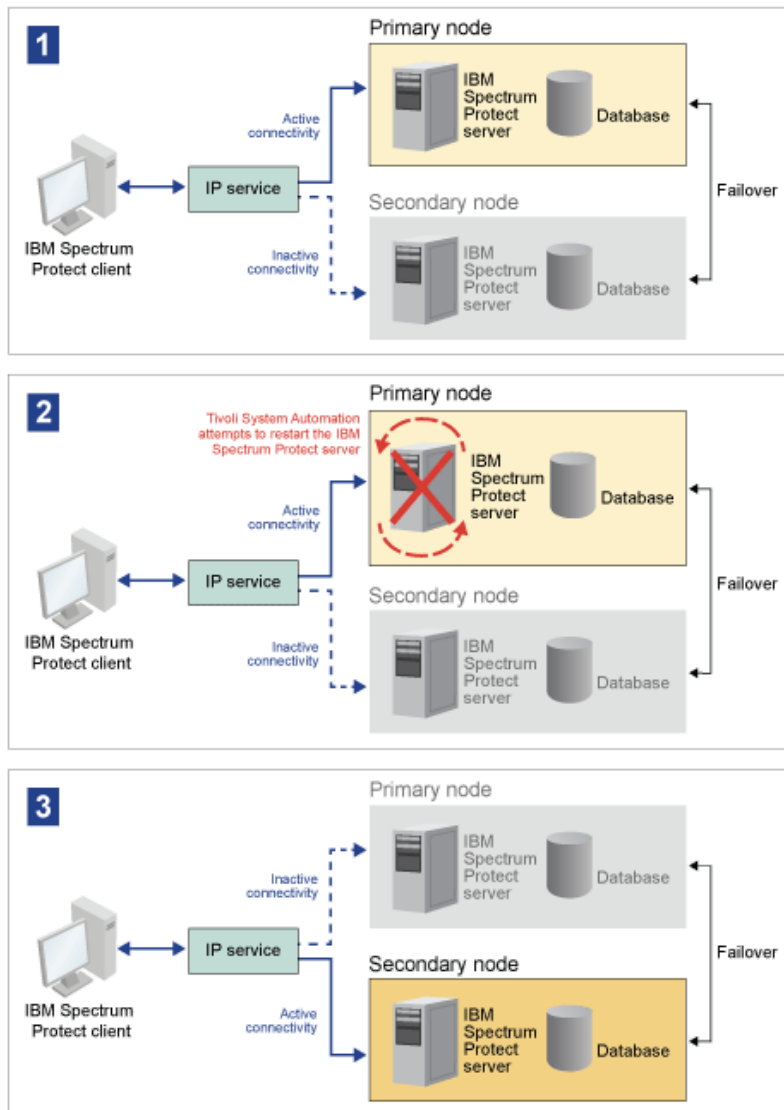
You can upgrade a server that is configured with System Automation for Multiplatforms.

Linux

Overview of a two-node IBM Spectrum Protect cluster using System Automation for Multiplatforms

Use the System Automation for Multiplatforms cluster for higher server and database availability during a failure. By using the System Automation for Multiplatforms failover function, server components such as the database can automatically recover from a failure.

The IBM Spectrum Protect™ server and the DB2® database are the underlying server components for this two-node cluster. The server is the core component. It is responsible for client and server activity. The DB2 database is an internal component, which is installed as part of the server. The server controls all database activity such as startup and shutdown. When the server detects a server or database component failure, it tries to restart the database. If the restart fails, the server and database are automatically shut down on the primary node and System Automation for Multiplatforms automatically starts these components on the secondary node. Because the IBM Spectrum Protect functions are restored immediately, server and database availability is higher. Figure 1. The failover function. The server and database components fail on the primary node. System Automation for Multiplatforms starts these components on the secondary node.




The server and the database include the following log directories, which are used for storage:

- IBM Spectrum Protect instance directory
- Active log directory
- Archive log directory
- Database directory

The two nodes in this System Automation for Multiplatforms cluster are configured to access highly available shared storage that protects the data. For example, a two-node topology includes a primary node and a secondary node. These nodes are on separate physical systems but they can access the same data by using the shared storage array.

Figure 2. Multiple IBM Spectrum Protect server instances on separate nodes. These server instances are on separate physical systems. These instances can access the highly available shared storage.



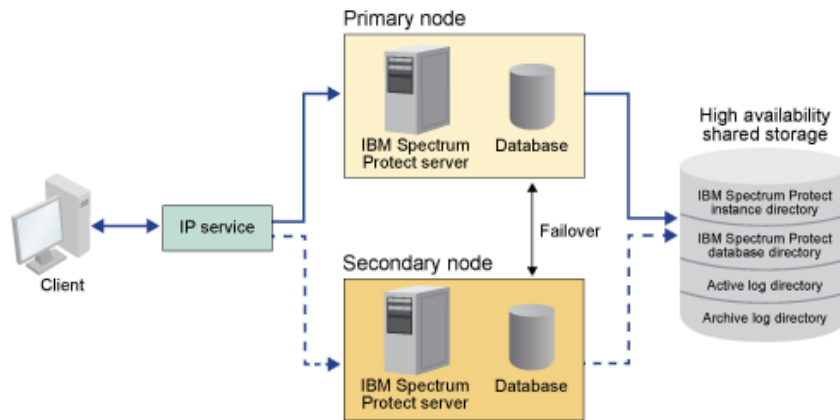
- Linux Two-node shared disk topology
 This cluster uses a two-node shared disk topology. It includes a primary and secondary node. The primary node hosts the IBM Spectrum Protect server, database, IBM Spectrum Protect instance, and the data. The secondary node is where the IBM Spectrum Protect resources are moved to if a failure occurs.
- Linux System Automation for Multiplatforms resource groups
 Use System Automation for Multiplatforms resources groups with defined automation policies to manage the IBM Spectrum Protect components for this cluster. The only exception is the database server instance resource that is managed by the IBM Spectrum Protect server.

Linux

Two-node shared disk topology

This cluster uses a two-node shared disk topology. It includes a primary and secondary node. The primary node hosts the IBM Spectrum Protect™ server, database, IBM Spectrum Protect instance, and the data. The secondary node is where the IBM Spectrum Protect resources are moved to if a failure occurs.

The two nodes in this cluster are connected to each other over a single public network and wired to a *shared disk storage* system, which is always available. *Shared disk storage* is where one or more disks are available to both the primary and secondary nodes. These disks are only mounted to one node, the primary node, at any one time. One node can input and output data to the shared storage disks. The following illustration shows a two-node shared topology where automatic failover to the secondary node occurs in the instance of a failure.



Linux

System Automation for Multiplatforms resource groups

Use System Automation for Multiplatforms resources groups with defined automation policies to manage the IBM Spectrum Protect™ components for this cluster. The only exception is the database server instance resource that is managed by the IBM Spectrum Protect server.

The shared file systems and IBM Spectrum Protect components are defined as resources. Multiple resources make up a resource group. Each resource in a resource group has a resource type. Each IBM Spectrum Protect instance in a cluster includes one resource group. During planned outages, resource groups can be manually moved from the primary node to the secondary node.

The IBM Spectrum Protect resource group includes the following resources. The name of the IBM Spectrum Protect resource group is SA-tsm-inst1-rg, where inst1 is the instance name. The following resources are used for different but mandatory functions in this cluster.

Service IP

The Service IP resource is used for communication. It is called tsm-inst1-ip-rs, where inst1 is the instance name. Service IP is managed by System Automation for Multiplatforms. This IP is available on the node where the IBM Spectrum Protect server is running. You must create the Service IP logical interface on the same physical interface as the public network interface.

Shared disk storage resource

A *shared disk storage* resource is a physical storage device on the IBM Spectrum Protect server where IBM Spectrum Protect and DB2® application data is stored. You must create the following disk storage resources:

- Instance directory - tsm-inst1-instdir-ag
- DB2 directory - tsm-inst1-db2dir-ag
- Active log directory - tsm-inst1-actlog-ag
- Archive log directory - tsm-inst1-archlog-ag

Shared disk storage for storage pools

The storage pool resource includes physical storage devices on the IBM Spectrum Protect server where client data is stored.

Volume group resources

If you decide to configure your storage by using volume groups, a volume group resource is available for the preceding *shared disk storage* resources. Volume group resources are automatically created by System Automation for Multiplatforms.

Application resources for the IBM Spectrum Protect server instance

The IBM Spectrum Protect server instance resource is the server resource that manages the IBM Spectrum Protect application. This resource is managed by System Automation for Multiplatforms control scripts.

Table 1. Tasks that are completed by the System Automation for Multiplatforms control scripts

Tasks	Description	Sample commands
Start	Starts the IBM Spectrum Protect server instance.	The <code>/opt/tivoli/tsm/server/bin/rc.dsmserv -u db2inst1 -i /tsminst1</code> command starts the server instance with the db2inst1 user in the /tsminst1 directory.
Stop	Stops the IBM Spectrum Protect server instance.	<code>kill -s SIGURG 345</code> where 345 is the <i>PID</i> . The <i>PID</i> can be found in the /tsminst1/dsmserv.v6lock file.
Monitor	Checks whether the /tsminst1/dsmserv.v6lock file exists. It uses the <i>PID</i> to check whether the process is running.	<code>ps -ef grep 345</code> where 345 is the <i>PID</i> .

- **Linux** Resource group dependencies
Resource group dependencies are automatically created to control the order in which resources are started. These dependencies also control which resources must be restarted or shut down when the specific resource that these resources depend on fails.

Linux

Setting up an IBM Spectrum Protect cluster with System Automation for Multiplatforms

You must set up the IBM Spectrum Protect™ cluster to use System Automation for Multiplatforms.

Procedure

1. Install and configure the IBM Spectrum Protect components on the primary and secondary nodes.
2. Install System Automation for Multiplatforms on the primary and secondary nodes.
3. Configure the storage resources.
4. Depending on the IBM Spectrum Protect version that is installed on the server, you might have to upgrade the IBM Spectrum Protect server for the System Automation for Multiplatforms cluster.
5. Optional: You can set the *FILE_EXIT* variable in the `tsmservctrl` cluster script to route the System Automation for Multiplatforms event data to the IBM Spectrum Protect server FILEEXIT file.
For example, edit the `tsmservctrl` cluster script in the `<server_install_directory>/tsam/controls` directory and add the following line:

```
FILE_EXIT="fileexittmp"
```

Linux

Prerequisites for configuring a Linux clustered environment with System Automation for Multiplatforms

Before you install and configure IBM Spectrum Protect™ in a clustered environment with System Automation for Multiplatforms, you must check the prerequisites.

Complete the following steps:

- Plan the installation of the IBM Spectrum Protect server. For more information, see *Installing and upgrading the server*.
- Prepare for the System Automation for Multiplatforms installation. For instructions, go to the System Automation for Multiplatforms product documentation. In the *Installation and Configuration Guide*, search for *Preparing for installation*.

Related tasks:

Planning to install the IBM Spectrum Protect server

Linux

Installing and configuring IBM Spectrum Protect components on the primary and secondary nodes

You must install the IBM Spectrum Protect™ server and database components on the primary and secondary nodes in the cluster. Then, configure the primary node first followed by the secondary node.

- **Linux** Installing IBM Spectrum Protect server components
After you verify the prerequisites, you must install the required IBM Spectrum Protect components.
- **Linux** Configuring the primary node
To set up the two-node topology, configure the IBM Spectrum Protect components on both nodes. First, you must configure the IBM Spectrum Protect instance on the primary node.
- **Linux** Configuring the secondary node
After you configure the primary node, you must configure the secondary node so that System Automation for Multiplatforms can move the IBM Spectrum Protect server components to the secondary node if the server fails on the primary node.

Linux

Installing IBM Spectrum Protect server components

After you verify the prerequisites, you must install the required IBM Spectrum Protect™ components.

Procedure

Install the IBM Spectrum Protect server on the primary and secondary nodes.

Related tasks:

Installing the IBM Spectrum Protect server components

Linux

Configuring the primary node

To set up the two-node topology, configure the IBM Spectrum Protect™ components on both nodes. First, you must configure the IBM Spectrum Protect instance on the primary node.

Before you begin

- Verify that the IBM Spectrum Protect instance owner has the same user and group ID for all of the nodes in the cluster domain.
- Verify that the IBM Spectrum Protect instance owner has the same password for all of the cluster nodes.

Procedure

1. For detailed instructions about creating the directories and the user ID for the server instance, see Linux: Creating the user ID and directories for the server instance.
2. Verify that the IBM Spectrum Protect server, DB2® instance, the active and archive log directories, and the mirror log directory, if applicable, are shared.
3. Define the mount points by adding entries to the `/etc/fstab` file.

When you add mount points on the cluster nodes, use the `noauto` option to prevent the mount points from being automatically mounted on more than one node in the cluster.

4. Set the following permissions on each of the mount points:
 - 755. For example, the following command sets the 755 permission on the `/tsminst1` mount point:

```
chmod -R 755 /tsminst1
```

- IBM Spectrum Protect server instance owner. For example, the following command sets the permissions for the instance owner:

```
chown -R tsminst1 /tsminst1
```

- IBM Spectrum Protect server group that the instance owner belongs to. For example, the following command sets the permissions for the instance owner's group:

```
chgrp tsmsrv_1_group /tsminst1
```

5. Mount the shared resources.
6. Log in to the primary node by using the instance user ID. Change to the instance directory and start the IBM Spectrum Protect server instance on the primary node by using the DSMSEV utility. For example, the following command starts the server for normal operation:

```
/opt/tivoli/tsm/server/bin/dsmsevr
```

7. Verify that the IBM Spectrum Protect components are started without any errors.
8. Halt the IBM Spectrum Protect server.
9. Log in with the root user ID and unmount the shared drives.

Linux

Configuring the secondary node

After you configure the primary node, you must configure the secondary node so that System Automation for Multiplatforms can move the IBM Spectrum Protect™ server components to the secondary node if the server fails on the primary node.

Procedure

1. To create the directories and the user ID for the server instance manually, follow the instructions in [Creating the user ID and directories for the server instance](#). Ensure that the same directory names are used on the primary and secondary nodes.
2. Define the mount points by adding entries in the `/etc/fstab` file.

When you add mount points on the cluster nodes, use the `noauto` option. This option prevents the mount points from being automatically mounted on more than one node in the cluster.

Ensure that the disk UIDs for each mount correspond to the disk UIDs on the primary node.

3. Set the following permissions on each of the mount points:
 - 755. For example, the following command sets the 755 permission on the `/tsminst1` mount point:

```
chmod -R 755 /tsminst1
```

- IBM Spectrum Protect server instance owner. For example, the following command sets the permissions for the instance owner:

```
chown -R tsminst1 /tsminst1
```

- IBM Spectrum Protect server group that the instance owner belongs to. For example, the following command sets the permissions for the instance owner's group:

```
chgrp tsmsrv_1_group /tsminst1
```

4. Mount the shared drives.
5. Create the IBM Spectrum Protect server instance by issuing the `db2icrt` command. For instructions, see [Creating the server instance](#).
Remember: You do not have to create a new server options file because the secondary node uses the `dsmsevr.opt` file from the primary node.
6. Log on to the secondary node by using the instance user ID. Catalog the database by issuing the `catalog db` command. For example, the following command catalogs the `tsmdb1` database:

```
db2 catalog db tsmdb1
```

7. Prepare the database for backup. For instructions, see [Preparing the database manager for database backup](#).
8. Change to the instance directory and start the IBM Spectrum Protect server by using the DSMSEV utility. For example, the following command starts the server for normal operation:

```
/opt/tivoli/tsm/server/bin/dsmsevr
```

9. Verify that the IBM Spectrum Protect components are starting without any errors.
10. Halt the IBM Spectrum Protect server and unmount the shared directories.

Installing System Automation for Multiplatforms on the primary and secondary nodes

After you install and configure IBM Spectrum Protect™ on the primary and secondary nodes in the cluster, you must install and configure System Automation for Multiplatforms on these nodes. Then, you must activate these nodes for the domain, configure the resources, and activate the base policy. Finally, you must add the mount points to the IBM Spectrum Protect directories.

- Linux** Creating the label for the mount points
 Create a label for each mount point on the primary and secondary nodes in the cluster.
- Linux** Installing and configuring System Automation for Multiplatforms
 You can integrate an IBM Spectrum Protect server with IBM Tivoli System Automation for Multiplatforms in a clustered environment. By using the System Automation for Multiplatforms failover function, you can help to ensure that IBM Spectrum Protect server components automatically recover from failures.
- Linux** Preparing to activate the cluster nodes for the domain
 After you install System Automation for Multiplatforms on the primary and secondary nodes in the cluster, you must prepare these nodes so that you can activate the cluster and start the cluster domain.
- Linux** Configuring volume group resources
 If you created volume groups for your cluster, you must configure these resources. System Automation for Multiplatforms automatically finds and defines the shared disk volume resources.
- Linux** Configuring resources that are not in a volume group
 If you created your *shared disk storage* resources by using ext2, ext3, or reiserfs resource types in one of the nodes in the cluster, then you must configure these resources.
- Linux** Activating the base policy
 After you configure the resources, you must activate the policy on the primary and secondary nodes to create any remaining resources and the resource group.
- Linux** Adding mount points to the IBM Spectrum Protect directories
 Before you can start the cluster, you must add the mount points that you created for the IBM Spectrum Protect components.

Creating the label for the mount points

Create a label for each mount point on the primary and secondary nodes in the cluster.

Procedure

1. Create a label for each of the volumes that you created previously for the shared directory mount points by issuing the `e2label` command. For example, the following command creates the `/tsminst1` label that has a `/dev/tsmvg1/tsminst1LV` partition.

```
e2label /dev/tsmvg1/tsminst1LV /tsminst1
```

2. For each node in the cluster, create an `e2label` entry for the mount points that you created previously in the `/etc/fstab` file. For example, for the previous sample label, issue the following command:

```
LABEL=/tsminst1 /tsminst1 ext3 defaults 0 0
```

Installing and configuring System Automation for Multiplatforms

You can integrate an IBM Spectrum Protect™ server with IBM® Tivoli® System Automation for Multiplatforms in a clustered environment. By using the System Automation for Multiplatforms failover function, you can help to ensure that IBM Spectrum Protect server components automatically recover from failures.

Before you begin

Complete the following tasks:

1. Ensure that you understand the basic terms, concepts, and components that are associated with System Automation for Multiplatforms. For more information, see Components.
2. Obtain the System Automation for Multiplatforms base release from Passport Advantage® and download the latest maintenance release from Fix Central.
3. Install and configure System Automation for Multiplatforms by following the instructions in the *Installation and Configuration Guide*.

Procedure

1. To integrate System Automation for Multiplatforms with IBM Spectrum Protect, follow the instructions in technote 7039780. Set up at least one cluster with at least two nodes.
2. Verify the configuration to ensure that all database, log, storage, and instance directories are on shared disks that are configured for failover. Take the following actions:
 - a. Log on as the instance user.
 - b. Run the `/opt/tivoli/tsm/server/bin/dsmclustfs` script and review the output. Ensure that the file systems that are reported by the script are on shared disks.
3. Test the clustered environment by using the procedures in technote 7039780. Ensure that the failover capabilities are working as expected.

Related information:

[IBM Tivoli System Automation for Multiplatforms Version 4.1 product information](#)

Linux

Preparing to activate the cluster nodes for the domain

After you install System Automation for Multiplatforms on the primary and secondary nodes in the cluster, you must prepare these nodes so that you can activate the cluster and start the cluster domain.

Procedure

1. Prepare each node for the domain by issuing the `preprnode` command. Issue this command for all the cluster nodes in the domain. For example, the following command prepares the `HOST1.ibm.com` and `HOST2.ibm.com` nodes:

```
preprnode HOST1.ibm.com HOST2.ibm.com
```

2. Create a domain by issuing the `mkrpdomain` command. For example, the following command creates the `tsm_domain` for the `HOST1.ibm.com` and `HOST2.ibm.com` nodes:

```
mkrpdomain tsm_domain HOST1.ibm.com HOST2.ibm.com
```

3. Start the domain for each node by issuing `startdomain` command. For example, the following command starts the `tsm_domain`:

```
startdomain tsm_domain
```

Linux

Configuring volume group resources

If you created volume groups for your cluster, you must configure these resources. System Automation for Multiplatforms automatically finds and defines the shared disk volume resources.

Procedure

To configure the volume group resources for the shared IBM Spectrum Protect™ directories and mount points that you created previously, complete the following steps on the primary node.

1. Import the volume groups. For example, use the `vgimport X` command to import the `x` volume groups.
2. Activate the volume groups. For example, use the `vgchange -ay X` command to activate the `x` volume groups.
3. Mount the file system by issuing the `mount` command. The following example mounts the `x` file system.

```
mount X
```

- Restart the domain by issuing the `stoprpdomain` and `startrpdomain` commands. For example, the following commands restart the `tsm_domain`.

```
stoprpdomain tsm_domain
startrpdomain tsm_domain
```

- Unmount the file system by issuing the `umount` command. For example, use the `umount X` command to unmount the `X` file system.
- Deactivate the volume groups. For example, use the `vgchange -an X` command to deactivate the `X` volume groups.
- Verify that all of the IBM®.AgfileSystem storage resources are harvested by System Automation for Multiplatforms by issuing the following command:

```
lsrsrc -s "Name=='Resource_Name' && ResourceType=1" IBM.AgFileSystem
```

Linux

Configuring resources that are not in a volume group

If you created your *shared disk storage* resources by using `ext2`, `ext3`, or `reiserfs` resource types in one of the nodes in the cluster, then you must configure these resources.

Procedure

Complete the following steps on the primary node.

- Mount the file system by issuing the `mount` command. For example, the following command mounts the `X` file system.

```
mount X
```

- Restart the domain by issuing the `stoprpdomain` and `startrpdomain` commands. For example, the following command restarts the `tsm_domain`.

```
stoprpdomain tsm_domain
startrpdomain tsm_domain
```

- Unmount the file system by issuing the `umount` command. For example, the following command unmounts the `X` file system.

```
umount X
```

- Verify that all of the IBM®.AgfileSystem storage resources are harvested by System Automation for Multiplatforms by issuing the following command:

```
lsrsrc -s "Name=='Resource_Name' && ResourceType=1" IBM.AgFileSystem
```

For example, to verify the `tsmalog` resource, issue the following command:

```
lsrsrc -s "Name=='tsmalog' && ResourceType=1" IBM.AgFileSystem
Resource Persistent Attributes for IBM.AgFileSystem resource 1:
ResourceHandle= "0x2038 0xffff 0x6ad47197 0x256fc23d 0x9338a9950x263fa510"
Name            = "tsmalog"
ResourceType    = 1  <-----
MountPoint      = ""
DeviceName      = ""
Vfs             = "ext3"
AggregateResource = "0x3fff 0xffff 0x00000000 0x00000000 0x00000000 0x00000000"
ContainerResource = "0x2036 0xffff 0x6ad47197 0x256fc23d 0x9338a995 0x25ffaa28"
GhostDevice     = 0
ResourceId      = "360050768019c021d30000000000005da"
ProtectionMode  = 1
UserControl     = 0
SysMountPoint   = "/tsmalog"
Label           = "/tsmalog"
FSID            = "5792f887-8547-4c33-a519-9d0c50ab6882"
PreOnlineMethod = 0
ContainerResourceId = "360050768019c021d30000000000005da"
AutoMonitor     = 1
Options         = "defaults,noauto"
PreOfflineMethod = 0
ActivePeerDomain = "TSM_Domain"
```

```
NodeNameList = {"tsmlnode01.storage.tucson.ibm.com", "tsmlnode02.storage.tucson.ibm.com"}
```

Linux

Activating the base policy

After you configure the resources, you must activate the policy on the primary and secondary nodes to create any remaining resources and the resource group.

About this task

To activate the base policy, you must create the Service IP resource and IBM Spectrum Protect™ application resources for the IBM Spectrum Protect server instance. Then, you must create the resource group and the policies to manage the cluster.

Procedure

Complete the following steps on all nodes in the cluster:

1. Go to the `/opt/tivoli/tsm/server/bin/tsam/bin` directory.
2. Update the following variables in the `base_cluster_variables.sh` script:
 - `NODE1` specifies the host name for node 1 (primary node) in the cluster.
 - `NODE2` specifies the host name for node 2 (secondary node) in the cluster.
 - `IP_GATEWAY` specifies the gateway of the Service IP.
 - `SUBNET_MASK` specifies the subnet mask of the Service IP.
 - `NET_INT` specifies the network interface name of a specific node in the cluster. This name must be the same for all the nodes in the cluster.
3. Run the `configureHA.sh` configuration script by issuing the `./configureHA.sh` command on all of the nodes in the cluster. If the `configureHA.sh` script fails with the `-bash: ./configureHA.sh: /bin/bash^M: bad interpreter: No such file or directory` error, issue the `dos2unix` command on all of the scripts in the bin directory. For example, for each script run the following command:

```
dos2unix -o <filename>
```

4. Verify that the configuration is a success by verifying that the configuration script runs successfully.
5. Attention: Complete this step on the primary node only.
Run the setup script by issuing the `./setup.sh` command. For example, the following command runs the setup script on the `inst1` IBM Spectrum Protect server instance for the `dbinst1` instance user in the `/tsminst1` IBM Spectrum Protect server instance directory with `9.11.142.129` as the service IP.

```
./setup.sh inst1 dbinst1 /tsminst1 9.11.142.129
```

6. Verify that an IP resource group was created by running the following command:

```
lssam -V
```

7. Repeat step 3 for all of the IBM Spectrum Protect instances that you have in your IBM Spectrum Protect server environment.

Linux

Adding mount points to the IBM Spectrum Protect directories

Before you can start the cluster, you must add the mount points that you created for the IBM Spectrum Protect™ components.

Procedure

To add the shared disk mount points to the cluster resource group and bring the cluster online, complete the following steps:

1. Identify mount points for the following directories:
 - Instance
 - Database
 - Active log
 - Archive log

- o Storage pool
- 2. Add resources to each mount point:
 - a. Verify whether the `tsm-$INST_NAME-rg` resource group is online by issuing the `lssam` command.
 - b. If the `tsm-$INST_NAME-rg` resource group is online, take it offline by issuing the following command:


```
chrg -o offline tsm-$INST_NAME-rg
```
 - c. Add shared disk resources to each mount point by running the `./update_setup.sh` script. For example, the following command adds the `/tsminst1` mount point to the `inst1` IBM Spectrum Protect server instance.


```
./update_setup.sh inst1 /tsminst1
```
- 3. Bring the `tsm-$INST_NAME-rg` resource group online by issuing the following command:


```
chrg -o online tsm-$INST_NAME-rg
```
- 4. Connect to the server using the service gateway IP to verify that the configuration is correct.

Linux

Configuring storage resources

Use the System Automation for Multiplatforms user interface or command line to add or delete storage resources and to delete mount points that are no longer required. If you add a storage pool to the cluster, you must add it to the resource group. If you remove a storage pool from the cluster, you must also delete it from the resource group.

- **Linux** Adding a storage pool to a resource group
If your IBM Spectrum Protect configuration stores data on disks, you must add the shared disk mount point for the storage pool to the resource group.
- **Linux** Deleting a storage pool from a resource group
You can delete a storage pool that is no longer required. If a storage pool is removed from the IBM Spectrum Protect server instance, it must be deleted from the resource group.
- **Linux** Deleting a mount point from a resource group
You might want to delete a mount point that is no longer required.

Linux

Adding a storage pool to a resource group

If your IBM Spectrum Protect™ configuration stores data on disks, you must add the shared disk mount point for the storage pool to the resource group.

Procedure

To add the shared disk mount point for the storage pool to the resource group, complete the following steps:

1. Lock the resource group by issuing the `rgreq -o lock` command. For example, the following command locks the `Sample_Resourcegroup_X` resource group:

```
rgreq -o lock Sample_Resourcegroup_X
```

2. Move to the `/opt/tivoli/tsm/server/bin/tsam/bin` directory.
3. To add a storage pool resource to a resource group, run the `update_setup.sh` script by issuing the `./update_setup.sh` command. For example, the following command adds the `/inst1stg1` storage pool mount point to the `inst1` IBM Spectrum Protect server instance:

```
./update_setup.sh inst1 /inst1stg1
```

4. Unlock the resource group by issuing the `rgreq -o unlock` command. For example, the following command unlocks the `Sample_Resourcegroup_X` resource group:

```
rgreq -o unlock Sample_Resourcegroup_X
```

Linux

Deleting a storage pool from a resource group

You can delete a storage pool that is no longer required. If a storage pool is removed from the IBM Spectrum Protect™ server instance, it must be deleted from the resource group.

Procedure

To delete a storage pool, complete the following steps:

1. Lock the resource group by issuing the `rgreq -o lock` command. For example, the following command locks the `Sample_Resourcegroup_X` resource group.

```
rgreq -o lock Sample_Resourcegroup_X
```

2. Move to the bin directory by issuing the `cd` command.
3. To delete a storage pool resource from a resource group, run the `delete_mount.sh` script by issuing the `./delete_mount.sh` command. For example, the following command deletes the `/inst1stg1` mount point from the `inst1` IBM Spectrum Protect server instance.

```
./delete_mount.sh /inst1stg1 inst1
```

4. Unlock the resource group by issuing the `rgreq -o unlock` command. For example, the following command unlocks the `Sample_Resourcegroup_X` resource group.

```
rgreq -o unlock Sample_Resourcegroup_X
```

Linux

Deleting a mount point from a resource group

You might want to delete a mount point that is no longer required.

Procedure

To delete a mount point, complete the following steps:

1. Check whether the `tsm-$INST_NAME-rg` resource group is online by issuing the `lssam` command.
2. If the `tsm-$INST_NAME-rg` resource group is online, take it offline by issuing the following command:

```
chrg -o offline tsm-$INST_NAME-rg
```

3. Move to the bin directory by issuing the `cd` command.
4. To delete a mount point, run the `delete_mount.sh` script. For example, the following command deletes the `/tsminst1` mount point from the `inst1` IBM Spectrum Protect™ server instance resource group.

```
./delete_mount.sh /tsminst1 inst1
```

5. Bring the `tsm-$INST_NAME-rg` resource group online by issuing the following command:

```
chrg -o online tsm-$INST_NAME-rg
```

Linux

Upgrading a server that is configured with System Automation for Multiplatforms

You can upgrade a server that is configured with System Automation for Multiplatforms.

Procedure

To upgrade the server on each node in the cluster, log in to the server and complete the following steps. These steps start the upgrade on the primary node and then the latter part of this procedure upgrades the secondary node.

1. Stop the server resources by issuing the `chrg -o Offline` command. For example, the following command stops the resources in the `tsm-tsminst1-rg` resource group:

```
chrg -o Offline tsm-tsminst1-rg
```

2. Stop the System Automation for Multiplatforms domain by issuing the `stoprpdomain` command. For example, the following command stops the `tsm_domain`:

```
stoprpdomain tsm_domain
```

3. Mount the server mount points on the primary node.
4. To upgrade the server on the primary node, see *Upgrading IBM Spectrum Protect™*.
5. After the upgrade is finished, complete the post upgrade steps to verify that the upgrade is successful on the primary node.
6. Stop the server and unmount the server mount points on the primary node.
7. Mount the server mount points on the secondary node.
8. If you are upgrading a server from V6 to V7, complete the following steps:
 - a. Uninstall the server.

For instructions, see *Uninstalling the V6.3 server* (see the *Installation Guide*).

b. Install the server on the secondary node. Follow the instructions in *Linux: Installing the server components*.

9. To upgrade the server on the secondary node, see *Upgrading the server*.
10. After the upgrade is complete, complete the post upgrade steps to verify that the upgrade is successful on the secondary node.
11. Unmount the server mount points on the secondary node.
12. Start the System Automation for Multiplatforms domain by issuing the `startrpdomain` command. For example, the following command starts the `tsa_domain`:

```
startrpdomain tsa_domain
```

13. Start the server resources by issuing the `chrg -o Online` command. For example, the following command starts the resources in the `tsm-tsminst1-rg` resource group:

```
chrg -o Online tsm-tsminst1-rg
```

Windows

Configuring a Windows clustered environment

You can configure an IBM Spectrum Protect™ server for Windows in a Microsoft failover cluster environment. Windows cluster environments consist of components such as IBM Spectrum Protect servers, hardware, and software. When these components are connected to the same disk system, downtime is minimized.

Microsoft software helps configure, monitor, and control applications and hardware components that are deployed on a Windows cluster. The administrator uses the Microsoft Cluster Administrator interface and IBM Spectrum Protect to designate cluster arrangements and define the failover pattern.

IBM Spectrum Protect supports tape failover for a cluster environment by using a Fibre or SCSI connection. Although Microsoft failover clusters do not support the failover of tape devices, the failover configuration can be monitored through the Microsoft Cluster Administrator interface after it is set up through IBM Spectrum Protect.

- **Windows** Microsoft Failover Cluster environment overview
With a Microsoft Failover Cluster Manager, you can place IBM Spectrum Protect server cluster resources into a cluster group. The IBM Spectrum Protect cluster group has a network name, an IP address, one or more physical disks, a Tivoli® server, and an IBM Spectrum Protect server service.
- **Windows** Tape failover for nodes in a cluster
Groups in a cluster can be transferred to other nodes when the node that is hosting the groups fails.
- **Windows** Planning for a clustered environment
Configuration in a clustered environment takes planning to ensure the optimal performance of your system. Whether you configure your system to include clusters depends on your business needs.
- **Windows** Setting up IBM Spectrum Protect in a Microsoft Failover Cluster
You must ensure that your cluster is properly installed and configured before you install IBM Spectrum Protect.
- **Windows** Maintaining the clustered environment
After you set up your initial cluster or clusters, maintenance needs are minimal.

Windows

Microsoft Failover Cluster environment overview

With a Microsoft Failover Cluster Manager, you can place IBM Spectrum Protect™ server cluster resources into a cluster group. The IBM Spectrum Protect cluster group has a network name, an IP address, one or more physical disks, a Tivoli® server, and an IBM Spectrum Protect server service.

The IBM Spectrum Protect instance network name is independent of the name of the physical node on which the IBM Spectrum Protect cluster group runs. Clients connect to an IBM Spectrum Protect server by using the instance network name, rather than the Windows node name. The instance network name maps to a primary or backup node. The mapping depends on which node owns the cluster group. Any client that uses Windows Internet Name Service (WINS) or directory services to locate servers can automatically track the IBM Spectrum Protect clustered server as it moves between nodes. You can automatically track the clustered server without modifying or reconfiguring the client.


Each IBM Spectrum Protect cluster group has its own disk as part of a cluster resource group. IBM Spectrum Protect cluster groups cannot share data between the cluster groups. Each IBM Spectrum Protect server that is configured in a cluster group has its database, active logs, recovery logs, and set of storage pool volumes on a separate disk. This disk is owned by the cluster group where the server is configured.

Remember: Microsoft Failover Cluster Manager supports an IP address only as a resource. Hence, any IBM Spectrum Protect server that runs on a cluster must limit its supported communication method to just TCP/IP. Any client that does not use TCP/IP as a communication method is not able to reach the IBM Spectrum Protect cluster group if it fails over to the other cluster node.

The following example demonstrates the way that a Microsoft Failover Cluster Manager for an IBM Spectrum Protect cluster server works.

Assume that a clustered IBM Spectrum Protect server that is named JUPITER is running on Node Z and a clustered IBM Spectrum Protect server that is named SATURN is running on Node X. Clients connect to the IBM Spectrum Protect server JUPITER and the IBM Spectrum Protect server SATURN without knowing which node hosts their server.

Figure 1. Clustering with JUPITER as Node Z and SATURN as Node X



When one of the software or hardware resources fails, failover occurs. Resources such as applications, disks, and an IP address move from the failed node to the remaining node. The remaining node:

- Takes over the IBM Spectrum Protect cluster group
- Brings the disk resources, the network resources, and the DB2 resource online
- Restarts the IBM Spectrum Protect service
- Provides access to administrators and clients

If Node X fails, Node Z assumes the role of running SATURN. To a client, it is exactly as if Node X were turned off and immediately turned back on again. Clients experience the loss of all connections to SATURN and all active transactions are rolled back to the client. Clients must reconnect to SATURN after the connection is lost. The location of SATURN is not apparent to the client.

Tape failover for nodes in a cluster

Groups in a cluster can be transferred to other nodes when the node that is hosting the groups fails.

A node can host physical or logical units, referred to as resources. Administrators organize these cluster resources into functional units that are called groups and assign these groups to individual nodes. If a node fails, the server cluster transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called *failover*. The reverse process, *failback*, occurs when the failed node becomes active again and the groups that were failed over to the other nodes are transferred back to the original node.

- **Windows** Fiber tape failover
IBM Spectrum Protect can manage the failover of Fibre Channel direct-attached tape and library devices on a Microsoft Windows system in a clustered environment without extra hardware.

Windows

Planning for a clustered environment

Configuration in a clustered environment takes planning to ensure the optimal performance of your system. Whether you configure your system to include clusters depends on your business needs.

Plan for a cluster configuration that accommodates your environment. In addition to assuring the correct type of hardware and the applicable software, you must set up a failover pattern.

When a node fails or needs to be taken offline, which node or nodes in the cluster picks up the transaction processing? In a two-node cluster, there is little planning necessary. In a more complex arrangement, you want to consider how your transaction processing is best handled. A form of load balancing among your nodes needs to be accounted for so that you maintain peak performance. Another consideration is to ensure that your customers do not see any lag and little drop in productivity.

Microsoft Cluster Servers and Microsoft Failover Clusters require each IBM Spectrum Protect™ server instance to have a private set of disk resources. Although nodes can share disk resources, only one node can actively control a disk at a time.

Attention: Ensure that the same level of Windows (Windows 2012, Windows 2012 R2, and Windows 2016) is installed on all computers in the cluster.

Is one configuration better than the other? To determine your best installation, you need to look at the differences in performance and cost. Assume that you have an IBM Spectrum Protect server-dedicated cluster whose nodes have comparable power. During failover, the performance of a configuration might degrade because one node must manage both IBM Spectrum Protect cluster instances. If each node handles 100 clients in a normal operation, one node must handle 200 clients during a failure.

- **Windows** Cluster configuration worksheet
Record your answers to the following planning questions before you set up the cluster configuration.
- **Windows** Preparing Windows systems for a clustered environment
You can prepare a Microsoft Windows system to host an IBM Spectrum Protect clustered environment.
- **Windows** Configuring IBM Spectrum Protect in Microsoft Failover Cluster
The IBM Spectrum Protect cluster configuration procedure must be completed on the set of nodes that hosts an IBM Spectrum Protect cluster group.

Related information:

[IBM Spectrum Protect Supported Operating Systems](#)

Windows

Cluster configuration worksheet

Record your answers to the following planning questions before you set up the cluster configuration.

1. What type of cluster solution best fits your business needs?
2. What type of failover pattern do you need?

The use of tape failover support also affects the pattern.

3. Is tape failover support be needed?

Consider how tape devices are used by the IBM Spectrum Protect™ cluster instances. The way that tape devices are used by cluster instances can limit the number of nodes in the failover pattern to two.

4. What are the resources to be dedicated to IBM Spectrum Protect?

Resource type	Resource name
Cluster Resource Group	
Physical Disk Resources	
IP address	
Subnet Mask	
Network	
Network Name (server name)	
Nodes	
Tape Failover (optional): device name - both nodes	

Windows

Preparing Windows systems for a clustered environment

You can prepare a Microsoft Windows system to host an IBM Spectrum Protect™ clustered environment.

Before you begin

Complete the following steps:

1. Ensure that the Windows cluster service is installed.
2. Install the Cluster Management Command Line tool, cluster.exe, which is available from the operating system installation disk. On the Windows system, in PowerShell, run the following commands:

```
PS C:\> Import-Module ServerManager
PS C:\> Add-WindowsFeature RSAT-Clustering-CmdInterface
```

Alternatively, run the following commands:

```
PS C:\> Import-Module ServerManager
PS C:\> Install-WindowsFeature -Name RSAT-Clustering-CmdInterface
```

Procedure

Complete the following steps:

1. Plan to host the server instance, the database, the logs, and the disk storage on shared disk.
2. Identify the disk resources to be dedicated to IBM Spectrum Protect. If you plan to install more than one IBM Spectrum Protect server, assign a different set of disks to each server.
3. Ensure that you have an IP address and network name for each IBM Spectrum Protect server instance that you plan to configure. For a cluster that involves two IBM Spectrum Protect cluster instances, you must have two network names.
4. Ensure that you have an IP address and network name for the cluster.
5. In the Windows cluster manager, create a cluster resource group and move disk resources to it. Each IBM Spectrum Protect server instance requires a cluster resource group. Initially, the group must contain only disk resources. You might choose to rename an existing resource group that contains only disk resources. IBM Spectrum Protect will be installed to a local disk on each node in the cluster.
6. Determine the disk to be used on each node. Plan to use the same drive letter on each system.

Windows

Configuring IBM Spectrum Protect in Microsoft Failover Cluster

The IBM Spectrum Protect™ cluster configuration procedure must be completed on the set of nodes that hosts an IBM Spectrum Protect cluster group.

Steps for the procedure vary depending upon which node you are currently configuring. When you configure the primary node in the set, the IBM Spectrum Protect server instance is created and configured. When you configure the remaining nodes in the set, each node is updated by using a specific method. The way the node is updated allows it to host the IBM Spectrum Protect server instance that is created on the primary node. An IBM Spectrum Protect server must be installed and configured on the first node in the set before you configure the remaining nodes in the set. Violating this requirement causes the configuration to fail.

Ensure that you completely configure one IBM Spectrum Protect cluster group before you move on to the next when you are configuring multiple IBM Spectrum Protect cluster groups. Because you are dealing with separate IP addresses and network names for each IBM Spectrum Protect cluster group, you lessen the possibility of mistakes by configuring each cluster group separately.

Windows

Setting up IBM Spectrum Protect in a Microsoft Failover Cluster

You must ensure that your cluster is properly installed and configured before you install IBM Spectrum Protect™.

Procedure

To configure IBM Spectrum Protect in a Microsoft Failover Cluster, complete the following steps:

1. Ensure that the Windows operating system is installed on all computers that are part of the cluster. For the most current information about supported Windows operating systems, see technote 1243309.
2. Log on with the domain user ID. The domain user must be in the same domain as the IBM Spectrum Protect server.
3. Ensure that the failover cluster is installed and configured for all the computers in the cluster.
If you plan to install the IBM Spectrum Protect server on the Windows Server 2012 operating system, install the failover cluster automation server and the failover cluster command interface first. To install these components, issue the following commands from Windows 2.0 PowerShell:

```
Install-WindowsFeature -Name RSAT-Clustering-AutomationServer  
Install-WindowsFeature -Name RSAT-Clustering-CmdInterface
```

4. Verify that each node and shared disk in the cluster is operational.
 5. Ensure that the shared tape devices are operational if IBM Spectrum Protect tape failover support is being used.
- **Windows** Preparing a Microsoft Failover Cluster group for a basic virtual server
Each IBM Spectrum Protect server instance requires a cluster resource group.
 - **Windows** Installing IBM Spectrum Protect in a Microsoft Failover Cluster
Install the IBM Spectrum Protect server on every node in the cluster that hosts an IBM Spectrum Protect clustered server.
 - **Windows** Initializing the IBM Spectrum Protect server for a Microsoft Failover Cluster on the primary node
After you install IBM Spectrum Protect on the nodes in the cluster, you must initialize the server on the primary node.
 - **Windows** Verifying the configuration of IBM Spectrum Protect in a Microsoft Failover Cluster
When you finish configuring IBM Spectrum Protect in a Microsoft Failover Cluster, you can review the Failover Cluster Manager summary window. Verify that clustering is completed successfully and the IBM Spectrum Protect server is started.
 - **Windows** Completing a failover test for your cluster
After you complete the cluster configuration, run a failover test to ensure that the nodes are working properly.

Windows

Preparing a Microsoft Failover Cluster group for a basic virtual server

Each IBM Spectrum Protect™ server instance requires a cluster resource group.

Before you begin

Use the Failover Cluster Manager program on the computer that owns the shared disk or tape resource to prepare your resource group. Initially, the group must contain only disk resources. You can create a group and move disk resources to it. You can also choose to rename an existing resource group that contains only disk resources.

As you construct your resource groups, consider the following items:

- Ensure that each resource group has a distinctive name. Do not change the names after the group is created because it can cause a corrupted configuration.

- Ensure that all nodes in the cluster are online.
- Ensure that the group is online and owned by the node where the initial server instance is installed.

Procedure

To prepare a resource group for cluster configuration, complete the following steps:

1. Open the Failover Cluster Manager program and expand the cluster. Right-click Roles and click Create Empty Role.
2. In the Roles pane, double-click New Role and change the role name to a meaningful name, such as TSMGROUP.
3. Right-click the resource group TSMGROUP and select Add storage.
4. In the Add storage area pane, select the shared volume or volumes for IBM Spectrum Protect and click OK. The resource group TSMGROUP, which contains the disk volumes that you added, is displayed.

Windows

Installing IBM Spectrum Protect in a Microsoft Failover Cluster

Install the IBM Spectrum Protect™ server on every node in the cluster that hosts an IBM Spectrum Protect clustered server.

Procedure

Complete the following steps for each node in your cluster to install the IBM Spectrum Protect server:

1. Log in with an administrator or domain user ID. The domain user must be a member of the Domain Administrators group.
2. Install the IBM Spectrum Protect server to a local disk on each node. Use the same local disk drive letter for each node.
3. Restart the system after the server installation completes.

Windows

Initializing the IBM Spectrum Protect server for a Microsoft Failover Cluster on the primary node

After you install IBM Spectrum Protect™ on the nodes in the cluster, you must initialize the server on the primary node.

Procedure

1. Ensure that all systems are restarted after the installation. Verify that all systems are running correctly.
2. Log in with an administrator or domain user ID. The domain user must be in the same domain as the IBM Spectrum Protect server.
3. Open the Failover Cluster Manager program and verify that the resources are online and owned by the primary node.
4. Begin the initialization procedure on the primary node in your cluster.
5. Temporarily enable the Microsoft Windows Server Message Block (SMBv1) protocol. Follow the instructions in How to detect, enable and disable SMBv1, SMBv2, and SMBv3 in Windows and Windows Server. This step is required to run the IBM Spectrum Protect configuration wizard.
6. From the Start menu, click All programs > IBM Spectrum Protect server > Configuration wizard.
7. Follow the wizard directions, clicking Next to step through the wizard. When you are prompted to enter the user ID, enter the name of the domain account to associate with the cluster.
8. If the initialization completed, click Done.
9. Disable the SMBv1 protocol. Follow the instructions in How to detect, enable and disable SMBv1, SMBv2, and SMBv3 in Windows and Windows Server.

Windows

Verifying the configuration of IBM Spectrum Protect in a Microsoft Failover Cluster

When you finish configuring IBM Spectrum Protect™ in a Microsoft Failover Cluster, you can review the Failover Cluster Manager summary window. Verify that clustering is completed successfully and the IBM Spectrum Protect server is started.

Procedure

To verify that the IBM Spectrum Protect server instance in a Microsoft Failover Cluster is created and configured correctly, complete the following steps:

1. From the Failover Cluster Manager, select the server instance. The network name that you configured is displayed in the Server Name pane.
2. In the Other Resources pane, confirm that the server instance and the IBM® DB2® server resource are displayed.
3. Right-click the IBM Spectrum Protect server instance and click Bring Online.

Windows

Completing a failover test for your cluster

After you complete the cluster configuration, run a failover test to ensure that the nodes are working properly.

Procedure

1. Open Failover Cluster Manager. Under Other Resources, right-click the IBM Spectrum Protect™ Instance(x) resource. Click Bring Online.
2. To test the failover, right-click the IBM Spectrum Protect cluster resource group and click Move.
3. Verify that the failover from the second node to the first node is completed successfully.

Windows

Maintaining the clustered environment

After you set up your initial cluster or clusters, maintenance needs are minimal.

Check your Windows Event log on a regular, if not daily, basis to monitor the activity of the nodes in the cluster. Use the log to check whether a node fails and needs maintenance.

The following list of topics describes situations that might affect the configuration or format of your cluster after it is operational.

- **Windows** Migrating an existing IBM Spectrum Protect server into a cluster
The reason for moving client data into a cluster is similar to the reason for adding a server to a cluster. You want to increase the availability and reliability of data to all your users. By having the server as part of the cluster, you provide an extra level of security by ensuring that no transactions are missed due to a failed server. The failover pattern that you establish prevents future failures.
- **Windows** Adding an IBM Spectrum Protect server with backup and restore
If your hardware resources are limited, you can add an existing IBM Spectrum Protect server to a cluster by using a backup and restore procedure.
- **Windows** Managing a virtual IBM Spectrum Protect server on a cluster
For most tasks, you can administer a virtual IBM Spectrum Protect server as you would a non-clustered server. To complete tasks such as starting and stopping the server or moving a resource group to another node to complete system maintenance, you must use the Microsoft Cluster Administrator interface.
- **Windows** Managing tape failover in a cluster
As part of your regular routine, check the event log to ensure that the configuration is operating properly. If a server fails, the error is logged. The log provides you with information to understand why the failure took place.
- **Windows** Troubleshooting with IBM Spectrum Protect cluster log
The IBM Spectrum Protect Cluster Resource DLL reports events and errors to the cluster log. The cluster log is a useful troubleshooting tool. When this log is enabled, it records the actions of each component of the Cluster service as the result of each action.

Windows

Migrating an existing IBM Spectrum Protect server into a cluster

The reason for moving client data into a cluster is similar to the reason for adding a server to a cluster. You want to increase the availability and reliability of data to all your users. By having the server as part of the cluster, you provide an extra level of security by ensuring that no transactions are missed due to a failed server. The failover pattern that you establish prevents future failures.

About this task

To migrate an existing IBM Spectrum Protect™ server into a cluster, you can either move the clients or complete a backup and restore procedure. The choice depends primarily on the availability and capacity of other IBM Spectrum Protect server computers in your site and your familiarity with the backup and restore procedure.

- **Windows** Moving the clients

If you move clients from a non-clustered IBM Spectrum Protect server computer to a clustered one, you can slowly move your users to the new system and not interrupt services. However, you must have the correct hardware to run two IBM Spectrum Protect servers simultaneously.

Related tasks:

Installing and upgrading the server

Windows

Adding an IBM Spectrum Protect server with backup and restore

If your hardware resources are limited, you can add an existing IBM Spectrum Protect™ server to a cluster by using a backup and restore procedure.

About this task

For example, suppose that you have no hardware other than the two server systems to be clustered. You plan to use the computer that is running the IBM Spectrum Protect server as a node. Complete this procedure to remove IBM Spectrum Protect from the computer and reinstall it in the cluster:

Procedure

1. Back up all disk storage pools to a copy storage pool.
2. Back up the database of the existing IBM Spectrum Protect server.
3. Perform the installation and configuration of the cluster.
4. Restore the database to the clustered IBM Spectrum Protect server.
5. Restore the disk storage pool volumes from the copy storage pool.
6. After you verify that all of your data is on the clustered server, delete the old server.

Windows

Managing a virtual IBM Spectrum Protect server on a cluster

For most tasks, you can administer a virtual IBM Spectrum Protect™ server as you would a non-clustered server. To complete tasks such as starting and stopping the server or moving a resource group to another node to complete system maintenance, you must use the Microsoft Cluster Administrator interface.

About this task

The Microsoft Cluster Administrator interface is available through the Administrative Tools program group. The interface is a detailed view of a virtual server configuration. The virtual server configuration includes details such as the physical Windows servers that are part of the cluster and their resources, network connections, and status. View the components of a virtual server configuration and start, stop, or fail back a virtual server by using this interface. Manage a virtual IBM Spectrum Protect server by using the Microsoft Cluster Administrator interface to avoid server failures and error messages. For example, if you use the Windows Service Control Manager to shut down the server, you might receive messages that the server failed.

You might want to move a virtual IBM Spectrum Protect server when the Windows server acts as the primary node and this server requires hardware or system maintenance. Use the Microsoft Cluster Administrator interface to move the management of the virtual IBM Spectrum Protect server to the secondary node until the maintenance is completed.

Windows

Managing tape failover in a cluster

As part of your regular routine, check the event log to ensure that the configuration is operating properly. If a server fails, the error is logged. The log provides you with information to understand why the failure took place.

Before you begin

To ensure that the server can identify or reset device names upon failover, set the SANDISCOVERY option to ON. By default, this option is set to OFF. For more information, see SANDISCOVERY.

About this task

Sometimes a node must rejoin the cluster, for example:

- When a node failed
- When a new Host Bus Adapter fiber card is added (equipment changes)

Procedure

Complete the following tasks in any order to ensure that a node can successfully join the cluster:

- Update, if necessary, the drive and library that use the IBM Spectrum Protect™ cluster tool.
- Take the IBM Spectrum Protect server offline until the failed node rejoins the cluster. This action helps ensure that the IBM Spectrum Protect server that is running on the other node is not affected.

Windows

Troubleshooting with IBM Spectrum Protect cluster log

The IBM Spectrum Protect™ Cluster Resource DLL reports events and errors to the cluster log. The cluster log is a useful troubleshooting tool. When this log is enabled, it records the actions of each component of the Cluster service as the result of each action.

In comparison with the Microsoft Windows Event Log, the cluster log is a complete record of cluster activity. The cluster log records the cluster service activity that is recorded in the event log. Although the event log can point you to a problem, the cluster log helps you resolve the problem.

The cluster log is enabled by default in Windows. Its output is printed as a log file in %SystemRoot%\Cluster. For more information, see the Windows online help documentation.

Configuring clients for applications, virtual machines, and systems

The server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.

- Adding clients
After you implement a data protection solution with IBM Spectrum Protect™, you can expand the solution by adding clients.
- Customizing policies
An organization's goals for how data is protected and retained are typically defined by corporate executives, legal advisors, or other people in lead roles. *Policies* are the means to align the operation of IBM Spectrum Protect with the data protection and retention goals of your organization.

Adding clients

After you implement a data protection solution with IBM Spectrum Protect™, you can expand the solution by adding clients.

About this task

The procedure describes basic steps for adding a client. For more specific instructions about configuring clients, see the documentation for the product that you install on the client node. You can have the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot

- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

Procedure

To add a client, complete the following steps:

1. Select the software to install on the client node and plan the installation. Follow the instructions in [Selecting the client software and planning the installation](#).
2. Specify how to back up and archive client data. Follow the instructions in [Specifying rules for backing up and archiving client data](#).
3. Specify when to back up and archive client data. Follow the instructions in [Scheduling backup and archive operations](#).
4. To allow the client to connect to the server, register the client. Follow the instructions in [Registering clients](#).
5. To start protecting a client node, install and configure the selected software on the client node. Follow the instructions in [Installing and configuring clients](#).

Selecting the client software and planning the installation

Different types of data require different types of protection. Identify the type of data that you must protect and select the appropriate software.

About this task

The preferred practice is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you install a product for which the client acceptor does not run schedules, you must follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

Procedure

Based on your goal, select the products to install and review the installation instructions.

Tip: If you install the client software now, you must also complete the client configuration tasks that are described in [Installing and configuring clients](#) before you can use the client.

Goal	Product and description	Installation instructions
Protect a file server or workstation	The backup-archive client backs up and archives files and directories from file servers and workstations to storage. You can also restore and retrieve backup versions and archived copies of files.	<ul style="list-style-type: none"> • Backup-archive client requirements • Install UNIX and Linux backup-archive clients • Installing the Windows client for the first time
Protect applications with snapshot backup and restore capabilities	IBM Spectrum Protect Snapshot protects data with integrated, application-aware snapshot backup and restore capabilities. You can protect data that is stored by IBM DB2® database software and SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications.	<ul style="list-style-type: none"> • Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux • Installing and upgrading IBM Spectrum Protect Snapshot for VMware • Installing and upgrading IBM Spectrum Protect Snapshot for Windows

Goal	Product and description	Installation instructions
Protect an email application on an IBM Domino® server	IBM Spectrum Protect for Mail: Data Protection for IBM® Domino automates data protection so that backups are completed without shutting down IBM Domino servers.	<ul style="list-style-type: none"> Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) Installation of Data Protection for IBM Domino on a Windows system (V7.1.0)
Protect an email application on a Microsoft Exchange server	IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server automates data protection so that backups are completed without shutting down Microsoft Exchange servers.	Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
Protect an IBM DB2 database	The application programming interface (API) of the backup-archive client can be used to back up DB2 data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)
Protect an IBM Informix® database	The API of the backup-archive client can be used to back up Informix data to the IBM Spectrum Protect server.	Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)
Protect a Microsoft SQL database	IBM Spectrum Protect for Databases: Data Protection for Microsoft SQL Server protects Microsoft SQL data.	Installing Data Protection for SQL Server on Windows Server Core
Protect an Oracle database	IBM Spectrum Protect for Databases: Data Protection for Oracle protects Oracle data.	Data Protection for Oracle installation
Protect an SAP environment	IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP provides protection that is customized for SAP environments. The product is designed to improve the availability of SAP database servers and reduce administration workload.	<ul style="list-style-type: none"> Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2 Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle
Protect a virtual machine	<p>IBM Spectrum Protect for Virtual Environments provides protection that is tailored for Microsoft Hyper-V and VMware virtual environments. You can use IBM Spectrum Protect for Virtual Environments to create incremental forever backups that are stored on a centralized server, create backup policies, and restore virtual machines or individual files.</p> <p>Alternatively, use the backup-archive client to back up and restore a full VMware or Microsoft Hyper-V virtual machine. You can also back up and restore files or directories from a VMware virtual machine.</p>	<ul style="list-style-type: none"> Installing Data Protection for Microsoft Hyper-V Installing and upgrading Data Protection for VMware Installing the IBM Spectrum Protect backup-archive clients (UNIX, Linux, and Windows)

Tip: To use the client for space management, you can install IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows.

Specifying rules for backing up and archiving client data

Before you add a client, ensure that appropriate rules are specified for backup and archive operations for the client data. During the client registration process, you assign the client node to a policy domain, which has the rules that control how and when client data is stored.

Before you begin

Determine how to proceed:

- If you are familiar with the policies that are configured for your solution and you know that they do not require changes, continue with Scheduling backup and archive operations.
- If you are not familiar with the policies, follow the steps in this procedure.

About this task

Policies affect how much data is stored over time, and how long data is retained and available for clients to restore. To meet objectives for data protection, you can update the default policy and create your own policies. A policy includes the following rules:

- How and when files are backed up and archived to server storage.
- The number of copies of a file and the length of time copies are kept in server storage.

During the client registration process, you assign a client to a *policy domain*. The policy for a specific client is determined by the rules in the policy domain to which the client is assigned. In the policy domain, the rules that are in effect are in the active *policy set*.

When a client backs up or archives a file, the file is bound to a management class in the active policy set of the policy domain. A *management class* is the key set of rules for managing client data. The backup and archive operations on the client use the settings in the default management class of the policy domain unless you further customize policy. A policy can be customized by defining more management classes and assigning their use through client options.

Client options can be specified in a local, editable file on the client system and in a client option set on the server. The options in the client option set on the server can override or add to the options in the local client option file.

Procedure

1. Review the policies that are configured for your solution by following the instructions in Viewing policies.
2. If you need to make minor changes to meet data retention requirements, follow the instructions in Editing policies.
3. Optional: If you need to create policy domains or make extensive changes to policies to meet data retention requirements, see Customizing policies.

Viewing policies

View policies to determine whether they must be edited to meet your requirements.

Procedure

1. To view the active policy set for a policy domain, complete the following steps:
 - a. On the Services page of the Operations Center, select a policy domain and click Details.
 - b. On the Summary page for the policy domain, click the Policy Sets tab.

Tip: To help ensure that you can recover data after a ransomware attack, apply the following guidelines:

 - Ensure that the value in the Backups column is a minimum of 2. The preferred value is 3, 4, or more.
 - Ensure that the value in the Keep Extra Backups column is a minimum of 14 days. The preferred value is 30 or more days.
 - Ensure that the value in the Keep Archives column is a minimum of 30 days.

If IBM Spectrum Protect™ for Space Management software is installed on the client, ensure that data is backed up before you migrate it. On the DEFINE MGMTCLASS or UPDATE MGMTCLASS command, specify MIGREQUIRESBKUP=YES. Then, follow the guidelines in the tip.
2. To view inactive policy sets for a policy domain, complete the following steps:
 - a. On the Policy Sets page, click the Configure toggle. You can now view and edit the policy sets that are inactive.
 - b. Scroll through the inactive policy sets by using the forward and back arrows. When you view an inactive policy set, the settings that differentiate the inactive policy set from the active policy set are highlighted.
 - c. Click the Configure toggle. The policy sets are no longer editable.

Editing policies

To change the rules that apply to a policy domain, edit the active policy set for the policy domain. You can also activate a different policy set for a domain.

Before you begin

Changes to policy can affect data retention. Ensure that you continue to back up data that is essential to your organization so that you can restore that data if a disaster occurs. Also, ensure that your system has sufficient storage space for planned backup operations.

About this task

You edit a policy set by changing one or more management classes within the policy set. If you edit the active policy set, the changes are not available to clients unless you reactivate the policy set. To make the edited policy set available to clients, activate the policy set.

Although you can define multiple policy sets for a policy domain, only one policy set can be active. When you activate a different policy set, it replaces the currently active policy set.

To learn about preferred practices for defining policies, see Customizing policies.

Procedure

1. On the Services page of the Operations Center, select a policy domain and click Details.
2. On the Summary page for the policy domain, click the Policy Sets tab.

The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.

3. Click the Configure toggle. The policy set is editable.
4. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.
5. Edit the policy set by completing any of the following actions:

Option	Description
Add a management class	<ol style="list-style-type: none">a. In the Policy Sets table, click +Management Class.b. To specify the rules for backing up and archiving data, complete the fields in the Add Management Class window.c. To make the management class the default management class, select the Make default check box.d. Click Add.
Delete a management class	In the Management Class column, click -. Tip: To delete the default management class, you must first assign a different management class as the default.
Make a management class the default management class	In the Default column for the management class, click the radio button. Tip: The default management class manages client files when another management class is not assigned to, or appropriate for managing, a file. To ensure that clients can always back up and archive files, choose a default management class that contains rules for both backing up and archiving files.
Modify a management class	To change the properties of a management class, update the fields in the table.

6. Click Save.
Attention: When you activate a new policy set, data might be lost. Data that is protected under one policy set might not be protected under another policy set. Therefore, before you activate a policy set, ensure that the differences between the previous policy set and the new policy set do not cause data to be lost.
7. Click Activate. A summary of the differences between the active policy set and the new policy set is displayed. Ensure that the changes in the new policy set are consistent with your data retention requirements by completing the following steps:
 - a. Review the differences between corresponding management classes in the two policy sets, and consider the consequences for client files. Client files that are bound to management classes in the active policy set will be bound to the management classes with the same names in the new policy set.
 - b. Identify management classes in the active policy set that do not have counterparts in the new policy set, and consider the consequences for client files. Client files that are bound to these management classes will be managed by the default management class in the new policy set.
 - c. If the changes to be implemented by the policy set are acceptable, select the I understand that these updates can cause data loss check box and click Activate.

Scheduling backup and archive operations

Before you register a new client with the server, ensure that a schedule is available to specify when backup and archive operations take place. During the registration process, you assign a schedule to the client.

Before you begin

Determine how to proceed:

- If you are familiar with the schedules that are configured for the solution and you know that they do not require modification, continue with Registering clients.
- If you are not familiar with the schedules or the schedules require modification, follow the steps in this procedure.


About this task

Typically, backup operations for all clients must be completed daily. Schedule client and server workloads to achieve the best performance for your storage environment. To avoid the overlap of client and server operations, consider scheduling client backup and archive operations so that they run at night. If client and server operations overlap or are not given enough time and resources to be processed, you might experience decreased system performance, failed operations, and other issues.

Procedure

1. Review available schedules by hovering over Clients on the Operations Center menu bar. Click Schedules.
2. Optional: Modify or create a schedule by completing the following steps:

Option	Description
Modify a schedule	<ol style="list-style-type: none">a. In the Schedules view, select the schedule and click Details.b. On the Schedule Details page, view details by clicking the blue arrows at the beginning of the rows.c. Modify the settings in the schedule, and click Save.
Create a schedule	In the Schedules view, click +Schedule and complete the steps to create a schedule.

3. Optional: To configure schedule settings that are not visible in the Operations Center, use a server command. For example, you might want to schedule a client operation that backs up a specific directory and assigns it to a management class other than the default.
 - a. On the Overview page of the Operations Center, hover over the settings icon  and click Command Builder.
 - b. Issue the DEFINE SCHEDULE command to create a schedule or the UPDATE SCHEDULE command to modify a schedule. For more information about the commands, see DEFINE SCHEDULE (Define a schedule for an administrative command) or UPDATE SCHEDULE (Update a client schedule).

Related tasks:

[Tuning the schedule for daily operations](#)

Registering clients

Register a client to ensure that the client can connect to the server, and the server can protect client data.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see technote 7048963.

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see Authenticating users by using an Active Directory database.

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the REGISTER NODE command and specify the USERID parameter:


```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see Register a node.

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:
 - a. On the Operations Center menu bar, click Clients.
 - b. In the Clients table, click +Client.
 - c. Complete the steps in the Add Client wizard:
 - i. Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the Enable check box.
 - ii. In the Configuration window, copy the TCPSEVERADDRESS, TCPPOINT, NODENAME, and DEDUPLICATION option values.
Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
 - iii. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
 - iv. Set how risks are displayed for the client by specifying the at-risk setting.
 - v. Click Add Client.

Related reference:

- [DECOMMISSION NODE \(Decommission a client node\)](#)
- [DECOMMISSION VM \(Decommission a virtual machine\)](#)
- [QUERY NODE \(Query nodes\)](#)
- [REMOVE REPLNODE \(Remove a client node from replication\)](#)

Installing and configuring clients

To start protecting a client node, you must install and configure the selected software.

Procedure

If you already installed the software, start at step 2.

1. Take one of the following actions:
 - o To install software on an application or client node, follow the instructions.

Software	Link to instructions
IBM Spectrum Protect™ backup-archive client	<ul style="list-style-type: none"> ▪ Install UNIX and Linux backup-archive clients ▪ Installing the Windows client for the first time <p>Tip: You can also update existing clients by using the Operations Center. For instructions, see Scheduling client updates.</p>
IBM Spectrum Protect for Databases	<ul style="list-style-type: none"> ▪ Data Protection for Oracle installation ▪ Installing Data Protection for SQL Server on Windows Server Core
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> ▪ Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) ▪ Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) ▪ Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none"> ▪ Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux ▪ Installing and upgrading IBM Spectrum Protect Snapshot for VMware ▪ Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none"> ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2® ▪ Installing IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle

- o To install software on a virtual machine client node, follow the instructions for the selected backup type.

Backup type	Link to instructions
-------------	----------------------

Backup type	Link to instructions
If you plan to create full VMware backups of virtual machines, install and configure the IBM Spectrum Protect backup-archive client.	<ul style="list-style-type: none"> ■ Install UNIX and Linux backup-archive clients ■ Installing the Windows client for the first time
If you plan to create incremental forever full backups of virtual machines, install and configure IBM Spectrum Protect for Virtual Environments and the backup-archive client on the same client node or on different client nodes.	<ul style="list-style-type: none"> ■ IBM Spectrum Protect for Virtual Environments online product documentation <p>Tip: You can obtain the software for IBM Spectrum Protect for Virtual Environments and the backup-archive client in the IBM Spectrum Protect for Virtual Environments installation package.</p>

2. To allow the client to connect to the server, add or update the values for the TCPSERVERADDRESS, TCPPORT, and NODENAME options in the client options file. Use the values that you recorded when you registered the client (Registering clients).
 - For clients that are installed on an AIX®, Linux, Mac OS X, or Oracle Solaris operating system, add the values to the client system-options file, dsm.sys.
 - For clients that are installed on a Windows operating system, add the values to the dsm.opt file.

By default, the options files are in the installation directory.
3. If you installed a backup-archive client on a Linux or Windows operating system, install the client management service on the client. Follow the instructions in Collecting diagnostic information with client management services.
4. Configure the client to run scheduled operations. Follow the instructions in Configuring the client to run scheduled operations.
5. Optional: Configure communications through a firewall. Follow the instructions in Configuring client/server communications through a firewall.
6. Run a test backup to verify that data is protected as you planned. For example, for a backup-archive client, complete the following steps:
 - a. On the Clients page of the Operations Center, select the client that you want to back up, and click Back Up.
 - b. Verify that the backup completes successfully and that there are no warning or error messages.
7. Monitor the results of the scheduled operations for the client in the Operations Center.

What to do next

If you need to change what is getting backed up from the client, follow the instructions in Modifying the scope of a client backup.

Configuring the client to run scheduled operations

You must configure and start a client scheduler on the client node. The client scheduler enables communication between the client and server so that scheduled operations can occur. For example, scheduled operations typically include backing up files from a client.

About this task

The preferred method is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations. The client acceptor manages the client scheduler so that the scheduler runs only when required:

- When it is time to query the server about the next scheduled operation
- When it is time to start the next scheduled operation

By using the client acceptor, you can reduce the number of background processes on the client and help to avoid memory retention problems.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Spectrum Protect™ for Databases, IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Mail, and IBM Spectrum Protect for Virtual Environments. If you installed a product for which the client acceptor does not run schedules, follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

If your business uses a third-party scheduling tool as standard practice, you can use that scheduling tool as an alternative to the client acceptor. Typically, third-party scheduling tools start client programs directly by using operating system commands. To

configure a third-party scheduling tool, see the product documentation.

Procedure

To configure and start the client scheduler by using the client acceptor, follow the instructions for the operating system that is installed on the client node:

AIX® and Oracle Solaris

- a. From the backup-archive client GUI, click Edit > Client Preferences.
- b. Click the Web Client tab.
- c. In the Managed Services Options field, click Schedule. If you also want the client acceptor to manage the web client, click the Both option.
- d. To ensure that the scheduler can start unattended, in the `dsm.sys` file, set the `passwordaccess` option to `generate`.
- e. To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

- f. Start the client acceptor by issuing the following command on the command line:

```
/usr/bin/dsmcad
```

- g. To enable the client acceptor to start automatically after a system restart, add the following entry to the system startup file (typically, `/etc/inittab`):

```
tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # Client Acceptor Daemon
```

Linux

- a. From the backup-archive client GUI, click Edit > Client Preferences.
- b. Click the Web Client tab.
- c. In the Managed Services Options field, click Schedule. If you also want the client acceptor to manage the web client, click the Both option.
- d. To ensure that the scheduler can start unattended, in the `dsm.sys` file, set the `passwordaccess` option to `generate`.
- e. To store the client node password, issue the following command and enter the client node password when prompted:

```
dsmc query sess
```

- f. Start the client acceptor by logging in with the root user ID and issuing the following command:

```
service dsmcad start
```

- g. To enable the client acceptor to start automatically after a system restart, add the service by issuing the following command at a shell prompt:

```
# chkconfig --add dsmcad
```

MAC OS X

- a. In the backup-archive client GUI, click Edit > Client Preferences.
- b. To ensure that the scheduler can start unattended, click Authorization, select Password Generate, and click Apply.
- c. To specify how services are managed, click Web Client, select Schedule, click Apply, and click OK.
- d. To ensure that the generated password is saved, restart the backup-archive client.
- e. Use the IBM Spectrum Protect Tools for Administrators application to start the client acceptor.

Windows

- a. In the backup-archive client GUI, click Utilities > Setup Wizard > Help me configure the Client Scheduler. Click Next.
- b. Read the information on the Scheduler Wizard page and click Next.
- c. On the Scheduler Task page, select Install a new or additional scheduler and click Next.
- d. On the Scheduler Name and Location page, specify a name for the client scheduler that you are adding. Then, select Use the Client Acceptor daemon (CAD) to manage the scheduler and click Next.
- e. Enter the name that you want to assign to this client acceptor. The default name is Client Acceptor. Click Next.
- f. Complete the configuration by stepping through the wizard.
- g. Update the client options file, `dsm.opt`, and set the `passwordaccess` option to `generate`.

h. To store the client node password, issue the following command at the command prompt:

```
dsmc query sess
```

Enter the client node password when prompted.

i. Start the client acceptor service from the Services Control page. For example, if you used the default name, start the Client Acceptor service. Do not start the scheduler service that you specified on the Scheduler Name and Location page. The scheduler service is started and stopped automatically by the client acceptor service as needed.

Configuring client/server communications through a firewall

If a client must communicate with a server through a firewall, you must enable client/server communications through the firewall.

Before you begin

If you used the Add Client wizard to register a client, find the option values in the client options file that you obtained during that process. You can use the values to specify ports.

About this task

Attention: Do not configure a firewall in a way that might cause termination of sessions that are in use by a server or storage agent. Termination of a valid session can cause unpredictable results. Processes and sessions might appear to stop due to input/output errors. To help exclude sessions from timeout restrictions, configure known ports for IBM Spectrum Protect™ components. Ensure that the KEEPALIVE server option remains set to the default value of YES. In this way, you can help to ensure that client/server communication is uninterrupted. For instructions about setting the KEEPALIVE server option, see KEEPALIVE.

Procedure

Open the following ports to allow access through the firewall:

TCP/IP port for the backup-archive client, command-line administrative client, and the client scheduler

Specify the port by using the `tcpport` option in the client options file. The `tcpport` option in the client options file must match the `TCPPORT` option in the server options file. The default value is 1500. If you decide to use a value other than the default, specify a number in the range 1024 - 32767.

HTTP port to enable communication between the web client and remote workstations

Specify the port for the remote workstation by setting the `httpport` option in the client options file of the remote workstation. The default value is 1581.

TCP/IP ports for the remote workstation

The default value of 0 (zero) causes two free port numbers to be randomly assigned to the remote workstation. If you do not want the port numbers to be randomly assigned, specify values by setting the `webports` option in the client options file of the remote workstation.

TCP/IP port for administrative sessions

Specify the port on which the server waits for requests for administrative client sessions. The value of the client `tcpadminport` option must match the value of the `TCPADMINPORT` server option. In this way, you can secure administrative sessions within a private network.

Scheduling client updates

Schedule the automatic installation of software updates for IBM Spectrum Protect™ backup-archive clients. This function is sometimes referred to as *client deployment*.

Before you begin

To schedule client updates by using the Operations Center, you must configure your environment to meet the following requirements:

Server requirements

IBM Spectrum Protect servers must meet the following requirements:

- IBM Spectrum Protect V8.1.3 or later must be installed on hub and spoke servers.
- Hub and spoke servers must have a high-level address and low-level address specified. You can configure these settings by using the SET SERVERHLADDRESS and SET SERVERLLADDRESS commands.
- The hub server must have a server password specified. You can configure this setting by using the SET SERVERPASSWORD command.
- The hub server must be defined to the spoke servers. This definition is not done automatically when spoke servers are added to the Operations Center. To define the hub server, issue the DEFINE SERVER command and use the second syntax option in the command documentation.

For example, issue the following command on each spoke server:

```
DEFINE SERVER hub_name SERVERPASSWORD=hub_pw HLA=hub_ip
LLA=hub_port SSL=NO SESSIONSECURITY=TRANSITIONAL
```

where the variables represent the following hub server settings: *hub_name* is the server name, *hub_pw* is the server password, *hub_ip* is the high-level address, and *hub_port* is the low-level address.

- The port that is specified by the RESTHTTPSPORT server option must be open to allow secure communication between the Operations Center and the hub server. The default port number is 8443.
- Spoke servers must have a directory-container storage pool or FILE storage pool available to store update packages. The Operations Center automatically selects a storage pool to use.

Client requirements

IBM Spectrum Protect backup-archive clients that you plan to update by using the Operations Center must meet the following requirements:

- The passwordaccess option must be set to generate.
- The autodeploy client option must be set to a value other than no. For more information about this option, see Autodeploy.
- Thirty-two-bit backup-archive clients are not supported. If a 32-bit backup-archive client is detected on a 64-bit operating system, the client is upgraded to the 64-bit version.
- The client scheduler must be running.
- The client system must be running and the client must have connected to the IBM Spectrum Protect server at least once.

Microsoft Windows clients must meet the following additional requirements:

- The client scheduler must be started as a Windows service and not from the command line. To minimize the chance of a restart, the scheduler service is shut down before the new client is installed, and restarted after the installation. If the scheduler is not run as a Windows service, a restart is required when the client is updated.
- The command-line version of the Windows registry utility (reg.exe) is required. This tool is generally installed as part of the operating system installation on supported Windows operating systems.

About this task

You can use the Operations Center to simultaneously update multiple clients at a scheduled time.

Update packages are automatically downloaded to the hub server, imported, and replicated to spoke servers. When an update schedule runs, files from the installation package are copied to the client system and the client is updated to the specified software version.

Restrictions:

- You can schedule only backup-archive client updates. Updates for other client types must be manually installed.
- The backup-archive client software cannot be updated by different IBM Spectrum Protect deployment managers at the same time.
- The Microsoft Windows cluster services environment is not supported.
- Do not schedule automatic client deployments to systems that have any of the following applications installed on them:
 - IBM Spectrum Protect for Virtual Environments
 - IBM Spectrum Protect for Databases
 - IBM Spectrum Protect for Mail
 - IBM Spectrum Protect for Enterprise Resource Planning
- To manage updates, the Operations Center creates policy objects, including device classes, storage pools, and domains, on the hub and spoke servers. The following naming convention is used for these objects: `IBM_DEPLOY_CLIENTS`. To avoid interfering with update operations, do not modify these objects.

- If you manually configured client deployment for an earlier server version, you must delete the policy objects that you defined before you schedule client updates by using the Operations Center.
- You can schedule updates to existing clients only. You cannot use the Operations Center to install a new client.

For information about manually installing backup-archive client software, see [Installing the IBM Spectrum Protect backup-archive clients in the IBM Spectrum Protect documentation](#).

For information about installing other IBM Spectrum Protect clients, see [Product suites and related products](#).

Procedure

1. On the Operations Center menu bar, click Updates > Clients. The Backup-Archive Client Updates page opens.
2. Use the information on the page to determine which release to install, click Schedule Update, and complete the steps in the wizard.

What to do next

To monitor, cancel, or reschedule updates, click Updates > Scheduled.

To diagnose and resolve issues, see [technote 2007749](#).

Related information:

SET SERVERHLADDRESS (Set the high-level address of a server)
SET SERVERLLADDRESS (Set the low-level address of a server)
SET SERVERPASSWORD (Set password for server)
DEFINE SERVER (Define a server for server-to-server communications)
RESTHTTSPORT

Customizing policies

An organization's goals for how data is protected and retained are typically defined by corporate executives, legal advisors, or other people in lead roles. *Policies* are the means to align the operation of IBM Spectrum Protect™ with the data protection and retention goals of your organization.

About this task

To automatically manage data protection and retention, you define policies, which are rules that you set on the server. Policies affect how much data is stored over time, and how long data is retained and available for clients to restore. Customize policies to meet the data protection goals for your organization.

You choose the policy that manages a client's data by assigning the client to a policy domain. Clients of different types have different retention requirements, and customizing and creating policies is typically necessary.

When a server is installed, by default it has one policy, in one policy domain. You can customize that policy and create your own policy.

- **Policy concepts**
The policy for a specific client is determined by the settings in the policy domain to which a client is added.
- **Customizing a policy**
You can customize existing policies to meet new or revised data retention requirements for your organization. Modifying a policy domain or copying an existing policy domain is a typical way to start customizing policy.
- **Creating a policy by copying an existing policy**
You can create new policies by copying an existing policy and then updating the parts that you want to change.
- **Creating a policy domain**
You might want to create a new policy domain for each type of client that is protected by the server. You might also want to divide responsibilities for clients among several administrators by giving them authority for specific policy domains.
- **Controlling client operations through client option sets**
You can use client option sets to centrally control the processing options that clients use for operations such as backup. Client option sets can help ensure that data is consistently protected according to your requirements. A client option set can override options in a local client options file, and can add options that might not be in a local client options file.

Policy concepts

The policy for a specific client is determined by the settings in the policy domain to which a client is added.

During the client registration process, you assign a client to a *policy domain*. The policy for each client is determined by the rules in the policy domain to which the client is assigned. In the policy domain, the rules that are in effect are in the active *policy set*.

When a client backs up or archives a file, the file is bound to a management class in the active policy set of the policy domain. A *management class* is the key set of rules for managing client data. The backup and archive operations on the client use the settings in the default management class of the policy domain unless you customize policy.

A policy can be customized by defining more management classes in the policy set, activating the policy set, and assigning the use of the new management classes through client options.

Client options can be specified in a local, editable file on the client system and in a client option set on the server. The options in the client option set on the server can override or add to the options in the local client option file.

The server uses the policy in management classes to manage files based on whether file versions are active or inactive. The most recent backup or archived copy of a file is the *active version*. Active versions are never deleted from server storage.

Backup versions other than the most recent version are called *inactive versions*. An active version of a file becomes inactive when one of the following events occur:

- The file is backed up again, creating a more recent version of the file in server storage.
- The file is deleted from storage on the client node and then an incremental backup operation runs. An *incremental backup*, the typical backup operation for a client, backs up only those files that changed since the last backup.

The settings in the management class that is bound to a file determine how long and how many inactive versions of the file are retained.

Expiration processing uses policies to determine when inactive versions are no longer needed, that is, when the versions are expired. The process of expiration on the server enforces policies that you define for data retention, and you must ensure that you schedule expiration to run regularly. For example, if you have a policy that requires only four versions of a file be kept, the fifth and oldest version is expired. During expiration processing, the server removes entries for expired versions from the database, which in effect deletes the versions from server storage.

- Retention and expiration of backup versions
Multiple versions of file backups are important because users can continually update files and might need to restore a file from different points in time. Policy settings control the backup versions that the server retains in server storage, and affects what users can restore.
- Policy activation after updates
When you make updates to policy, the updates do not go into effect until you activate the policy set that you updated.

Related information:

[Full incremental backup and partial incremental backup](#)

Retention and expiration of backup versions

Multiple versions of file backups are important because users can continually update files and might need to restore a file from different points in time. Policy settings control the backup versions that the server retains in server storage, and affects what users can restore.

You can specify the versions that the server retains in server storage with settings in the management class:

- Specify the number of days to keep backup versions.
You specify the number of days to keep backup versions with settings in the Operations Center:
 - Keep Extra Backups, which is how many days to keep inactive backup versions. The days are counted from the day that the version becomes inactive.

If you use commands, use the DEFINE COPYGROUP command with the RETEXTRA parameter.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 14 days. The preferred value is 30 or more days.
 - Keep Deleted Backups, which is how many days to keep the last backup version of a file that is deleted from the client file system.

If you use commands, use the DEFINE COPYGROUP command with the RETONLY parameter.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 30 days.

- Specify the number of versions to keep.

You specify the number of backup versions to keep with settings in the Operations Center:

- Backups, which is the number of versions to keep of a file that still exists on the client file system.

If you use commands, use the DEFINE COPYGROUP command with the VEREXISTS parameter.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 2. Preferred values are 3, 4, or more.

- Deleted Backups, which is the number of versions to keep of a file that is deleted from the client file system.

If you use commands, use the DEFINE COPYGROUP command with the VERDELETED parameter.

- Specify a combination of the number of versions and the days to keep them.

The settings interact to determine the backup versions that the server retains. Ensure that you understand which settings take precedence and what interactions can occur:

- When the number of inactive backup versions exceeds the number in the Backups and Deleted Backups settings, the oldest version expires and the server deletes the version from the database the next time expiration processing runs.
- The number of inactive versions that the server keeps is also affected by the Keep Extra Backups setting. Inactive versions expire when the number of days that they are inactive exceeds the value that is specified for retaining extra versions, even when the number of versions is not exceeded.

- File expiration and expiration processing

Files expire as they exceed the retention criteria that is specified in the policy. Expiration processing on the server removes expired files from the server database and the files are deleted from server storage.

- Example: Retention when a policy uses only time controls

The simplest way to manage data retention is to use only time-based policy controls. With only time-based controls in the policy, the file versions are retained based on the days since the versions become inactive.

- Example: Retention when a policy uses both version and time controls

Using both the version and the time controls in a policy gives you flexibility in managing data retention but also causes complexity. To understand the interactions among the controls, review example policies and their effects on the retention of one file's backup versions during one month.

- Interactions among policy settings

Time-based and version-based policy settings interact when used together in a management class for a policy. The frequency of client backups also affects the backup versions that are stored for a client.

File expiration and expiration processing

Files expire as they exceed the retention criteria that is specified in the policy. Expiration processing on the server removes expired files from the server database and the files are deleted from server storage.

Files expire under the following conditions:

- Users delete file spaces from client nodes
- Users expire files by using the EXPIRE command on the client
- A backup version of a file exceeds the criteria for backup retention (how long a file is kept and how many inactive versions of a file are kept)
- An archived file exceeds the time criteria for retention of archived files (how long archived copies are kept)
- A backup set exceeds the retention time that is specified for the backup set

The server deletes expired files from the server database only during expiration processing. After expired files are deleted from the database, the server can reuse the space in the storage pools that was occupied by expired files. Ensure that expiration processing runs periodically to allow the server to reuse space.

Restrictions on expiration processing

The use of some functions affects expiration processing.

Replication

If you are using dissimilar policies on the source and target servers, files that are marked for immediate expiration on the source replication server are not deleted until they are replicated to the target replication server. If you are not using dissimilar policies, files that are marked for immediate expiration on the source replication server are deleted immediately.

For the target replication server, if files are marked as expired, they are deleted when the target replication server runs expiration processing.

Event-based retention for archive data

An archive file is not eligible for expiration if there is a deletion hold on it. If a file is not held, it is handled according to existing expiration processing.

Related information:

[Expiration/deletion hold and release](#)

Example: Retention when a policy uses only time controls

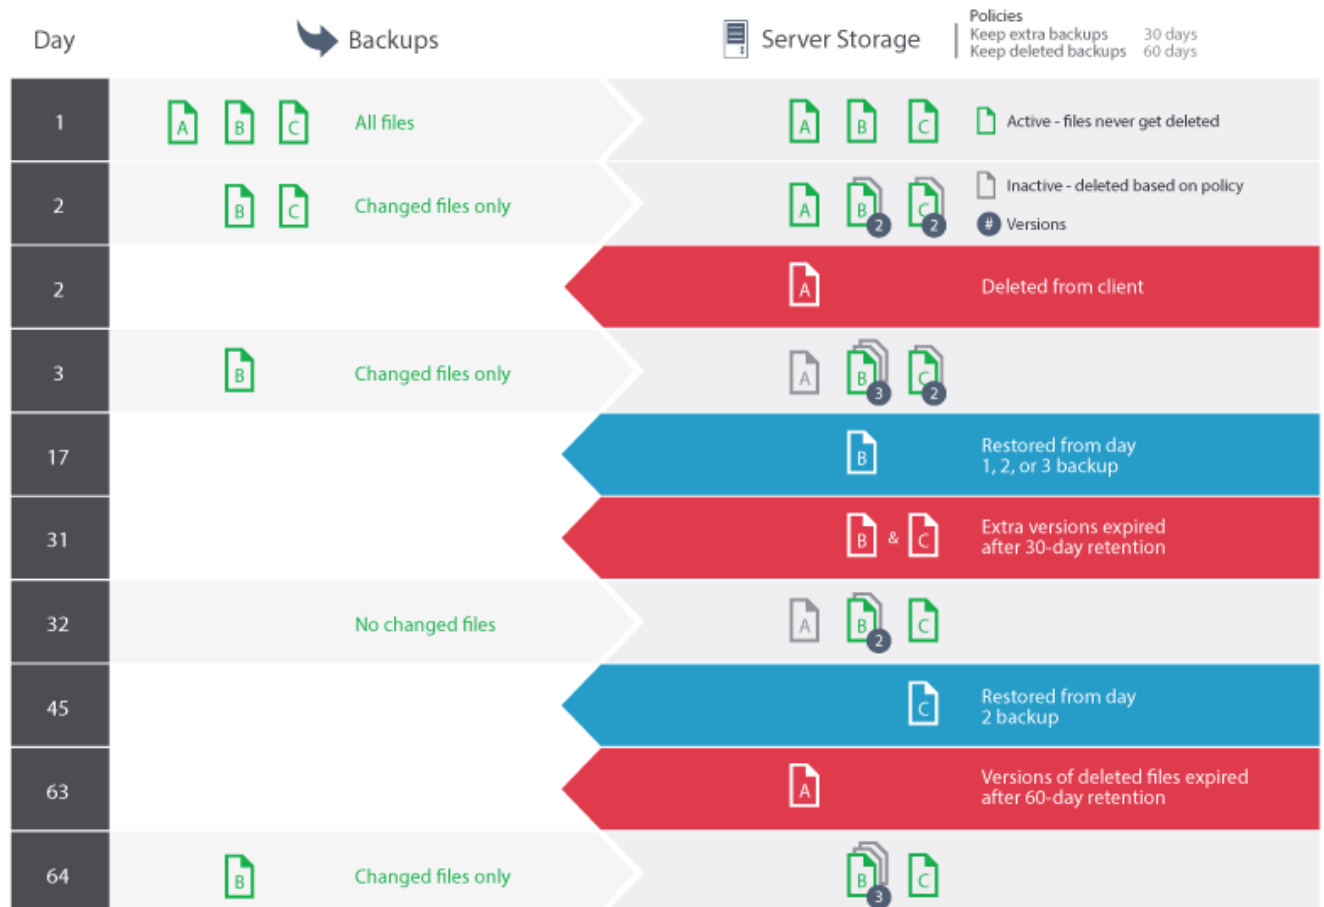
The simplest way to manage data retention is to use only time-based policy controls. With only time-based controls in the policy, the file versions are retained based on the days since the versions become inactive.

For a policy that is based only on time, you use the Keep Extra Backups and Keep Deleted Backups controls. This type of policy does not limit the number of versions of files. If clients back up frequently, ensure that server storage can handle the potential number of file versions.

The following figure shows how files from a client are handled by the server over time as the client runs a daily incremental backup operation.

In this example, the policy has the following characteristics:

- The latest version of a file is always retained, as long as the file still exists on the client system. The latest version is the active version. This characteristic is part of every policy on the server.
- Keep Extra Backups is set to 30 days. After a more recent backup is made, a file version becomes inactive and is kept in server storage for 30 days.
- Keep Deleted Backups is set to 60 days. When a file is deleted from the client system, all versions of the file in server storage become inactive. These inactive versions are kept for 60 days after the file versions become inactive.



Example: Retention when a policy uses both version and time controls

Using both the version and the time controls in a policy gives you flexibility in managing data retention but also causes complexity. To understand the interactions among the controls, review example policies and their effects on the retention of one file's backup versions during one month.

See Table 1 and Figure 1. A client backs up the file REPORT.TXT four times in one month, from 23 March to 23 April. The settings in the backup copy group of the management class to which REPORT.TXT is bound determine how the server treats these backup versions. Table 2 shows how different copy group settings can affect the versions, as of 24 April (one day after the file was last backed up).

Table 1. Status of REPORT.TXT backup versions as of 24 April

Version	Date created	Days since the version became inactive
Active	23 April	(not applicable)
Inactive 1	13 April	1 (since 23 April)
Inactive 2	31 March	11 (since 13 April)
Inactive 3	23 March	24 (since 31 March)

Figure 1. Active and inactive versions of REPORT.TXT




Table 2. Effects of the policy on retention of backup versions for REPORT.TXT as of 24 April

Backups	Deleted Backups	Keep Extra Backups	Keep Deleted Backups	Results
4 versions	2 versions	60 days	180 days	<p>Backups and Keep Extra Backups settings control the expiration of the versions. The version that is created on 23 March is retained until the file is backed up again (creating a fourth inactive version), or until that version is inactive for 60 days.</p> <p>If the user deletes the REPORT.TXT file from the client file system, the server notes the deletion at the next full incremental backup operation by the client. From that point, the Deleted Backups and Keep Deleted Backups settings also affect the retention. All versions are now inactive.</p> <p>Two of the four versions expire immediately (the 23 March and 31 March versions expire). The 13 April version expires when it is inactive for 60 days (on 23 June). The server keeps the last remaining inactive version, the 23 April version, for 180 days after it becomes inactive.</p>

Backups	Deleted Backups	Keep Extra Backups	Keep Deleted Backups	Results
No limit	2 versions	60 days	180 days	<p>Keep Extra Backups setting controls expiration of the versions. The inactive versions (other than the last remaining version) are expired when they are inactive for 60 days.</p> <p>If the user deletes the REPORT.TXT file from the client node, the server notes the deletion at the next full incremental backup operation by the client. From that point, the Deleted Backups and Keep Deleted Backups settings also affect the retention. All versions are now inactive.</p> <p>Two of the four versions expire immediately (the 23 March and 31 March versions expire) because only two versions are allowed. The 13 April version expires when it is inactive for 60 days (on 22 June). The server keeps the last remaining inactive version, the 23 April version, for 180 days after it becomes inactive.</p>
No limit	No limit	60 days	180 days	<p>Keep Extra Backups setting controls expiration of the versions. The server does not expire inactive versions based on the maximum number of backup copies. The inactive versions (other than the last remaining version) are expired when they are inactive for 60 days.</p> <p>If the user deletes the REPORT.TXT file from the client node, the server notes the deletion at the next full incremental backup operation by the client node. From that point, the Keep Deleted Backups setting also affects the retention. All versions are now inactive.</p> <p>Three of the four versions expire after each of them is inactive for 60 days. The server keeps the last remaining inactive version, the 23 April version, for 180 days after it becomes inactive.</p>
4 versions	2 versions	No limit	No limit	<p>Backups setting controls the expiration of the versions until a user deletes the file from the client node. The server does not expire inactive versions based on age.</p> <p>If the user deletes the REPORT.TXT file from the client node, the server notes the deletion at the next full incremental backup operation by the client node. From that point, the Deleted Backups setting controls expiration. All versions are now inactive.</p> <p>Two of the four versions expire immediately (the 23 March and 31 March versions expire) because only two versions are allowed. The server keeps the two remaining inactive versions indefinitely.</p>

Related information:

[Full incremental backup and partial incremental backup](#)

Interactions among policy settings

Time-based and version-based policy settings interact when used together in a management class for a policy. The frequency of client backups also affects the backup versions that are stored for a client.

For a client system that must back up twice a day, consider the effects of the following policy choices on a file that changes frequently:

- You set Keep Extra Backups to 30 days. You set Backups to No limit so that the policy does not limit the number of versions. After 30 days, the server might have 60 backup versions of the file, if the file changes between each of the two daily backup operations. The client can choose to restore any of the 60 versions from the past 30 days.
- You set Keep Extra Backups to No limit, and set Backups to 30 versions. If the file changes between each of the two daily backup operations, the server has 30 backup versions after 15 days. After 30 days, the server still has only 30 backup versions because of the limit on the number of versions. If the file continues to change between each of the two daily backup operations, the backup versions might be from as few as the last 15 days. The client can choose to restore one of the 30 versions, which might be no older than 15 days.

The examples show that if backup versions must be available for a specific number of days, the simplest way to implement that requirement is to use a time-based policy. Set Keep Extra Backups to the specific number of days and set Backups to No limit.

The effect of the No limit value in the policy settings varies according to what other policy controls are set:

Keep Extra Backups

If you specify No limit, inactive backup versions are deleted based on the Backups or Deleted Backups settings.

To enable client nodes to restore files to a specific point in time, set the Backups or Deleted Backups to No limit. Set Keep Extra Backups to the number of days that you expect clients might need versions of files available for possible point-in-time restoration. For example, to enable clients to restore files from a point in time 60 days in the past, set Keep Extra Backups to 60.

Keep Deleted Backups

If you specify No limit, the last version is retained forever unless a user or an administrator deletes the file from server storage.

Backups

Setting the value to No limit can require increased storage, but that value might be necessary to specify for some situations. For example, to enable client nodes to restore files to a specific point in time, set the value for Backups to No limit. By not setting a limit on versions, you ensure that the server retains versions according to the Keep Extra Backups setting.

Deleted Backups

Setting the value to No limit can require increased storage, but that value might be necessary to specify for some situations. For example, set the value for Deleted Backups to No limit to enable clients to restore files to a specific point in time. By not setting a limit on versions, you ensure that the server retains versions according to the Keep Extra Backups setting.

Cross-reference for Operations Center fields and server command parameters

The following table shows the Operations Center fields with the equivalent parameter to use with the `DEFINE COPYGROUP TYPE=BACKUP` command.

Field name in Operations Center views	Parameter to use with the <code>DEFINE COPYGROUP TYPE=BACKUP</code> command
Keep Extra Backups	RETEXTRA
Keep Deleted Backups	REONLY
Backups	VEREXISTS
Deleted Backups	VERDELETED

Policy activation after updates

When you make updates to policy, the updates do not go into effect until you activate the policy set that you updated.

Policy set activation puts into effect updates that you made. For example, the following types of updates go into effect after you activate the policy set:

- You define a new policy domain with a policy set and one or more management classes
- You add a management class to a policy set
- You change the backup retention settings in an existing management class

Policy set validation before activation

In the Operations Center, validation is not a separate step. If you are using commands, validation is an optional command that gives you an opportunity to preview the effect of activating a changed policy set. When you validate a policy set, the server reports on conditions that might cause problems if the policy set is activated.

Validation fails if the policy set does not contain a default management class. Validation results in warning messages if any of the conditions that are shown in Table 1 exist.

Table 1. Conditions that cause warnings during policy set validation

Condition	Reason for warning
The storage destinations that are specified for backup, archive, or migration operations are not defined storage pools.	A storage pool must exist before you can specify it as a destination.
A storage destination that is specified for backup, archive, or migration operations is a copy storage pool or an active-data pool.	The storage destination must be a primary storage pool.
The default management class does not contain backup or archive settings.	When the default management class does not contain backup or archive settings, any files that are bound to the default management class are not backed up or archived.
The current active policy set names a management class that is not defined in the policy set that is being validated.	<p>When you back up files that were bound to a management class that no longer exists in the active policy set, backup versions are rebound to the default management class.</p> <p>When the management class to which an archive copy is bound no longer exists and the default management class does not contain archive settings, the server uses the archive retention grace period to manage the retention of the archive copy.</p> <p>The archive retention grace period is set for a policy domain, and that setting is used only when no other policy setting is available to manage an archive copy.</p>
The current active policy set contains backup settings that are not defined in the policy set that is being validated.	<p>When a client backs up a file and the management class to which the file is bound no longer has backup settings, backup versions are managed by the default management class.</p> <p>If the default management class does not contain any backup settings, the server uses the backup retention grace period to manage file versions. However, the file is not backed up in the next backup operation.</p> <p>The backup retention grace period is set for a policy domain, and that setting is used only when no other policy setting is available to manage a backup version.</p>
A management class specifies that a backup version must exist before a file can be migrated from a client node, but the management class does not contain backup settings.	This warning applies only if you are using the IBM Spectrum Protect™ for Space Management product. The conflicts within the management class can cause problems for IBM Spectrum Protect for Space Management clients.

Policy set activation

When you activate a policy set, the server validates the contents of the policy set and copies the policy set to be the active policy set. To later change the contents of the active policy set, you must create or change another policy set and then activate that policy set.

Some updates to a policy have an immediate effect when it is activated, but some other updates do not:

- Updates to Keep Extra Backups and Keep Deleted Backups settings are immediately applied to data already in server storage, and to future backups.

If you use commands, these settings are the RETEXTRA and RETONLY parameters for the DEFINE COPYGROUP or UPDATE COPYGROUP commands.
- Updates to Backups and Deleted Backups settings do not take effect for client data until the clients complete the next backup operation.

If you use commands, these settings are the VEREXISTS and VERDELETED parameters for the DEFINE COPYGROUP or UPDATE COPYGROUP commands.

Restrictions for servers that use the feature for data retention protection

If the feature for data retention protection is active, more rules apply when you validate and activate a policy set. The feature for data retention protection is activated by using the SET ARCHIVERETENTIONPROTECTION command on a server that does not yet have any client data.

If data retention protection is active for a server, more rules must be satisfied before the policy is activated:

- If a management class exists in the active policy set, a management class with the same name must exist in the policy set that is being activated.
- All management classes in the policy set that is being activated must contain archive retention settings.
- If the active policy set includes archive retention settings in a management class, the policy set that is being activated must have archive retention values at least as large as the corresponding values in the active policy set.

If the server is a managed server in an enterprise configuration, the server might receive policy updates from the server that is the configuration manager. Policy updates that are received by the managed server from the configuration manager must also satisfy the preceding rules.

Related concepts:

[Enterprise configuration \(V7.1.1\)](#)

Related reference:

SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)

Customizing a policy

You can customize existing policies to meet new or revised data retention requirements for your organization. Modifying a policy domain or copying an existing policy domain is a typical way to start customizing policy.

About this task

The key policy settings are in management classes. In the management classes, you can control both the number of backup versions and the number of days that backup versions are kept in server storage. When you use both types of controls, the policy is more complex. By controlling only the number of days that backup versions are kept, you can more simply define how long backed-up data is kept.

Ensure that the default management class in a policy domain has appropriate settings for data retention for most or all of the clients that are assigned to the domain. The retention settings in the default management class are applied to data when client operations do not specify a management class.

You can work on updates to a policy and save the changes until a later time. When you are satisfied that the draft changes are ready, you can activate the updated policy set to put the changes into effect.

Procedure

1. On the Overview page of the Operations Center, click the Services menu.
2. Select the policy domain and click Details. Click Policy Sets.
3. Click the Configure toggle so that you can update the settings.
4. Customize the settings in the management class.
 - a. Make selections for backup services. For example, update the following items so that inactive backup versions for the clients are retained for 30 days:
 - Backups: No limit
 - Keep Extra Backups: 30 days
 - Deleted Backups: 1
 - Keep Deleted Backups: No limit
 - b. Optional: Make selections for archive services. For example, change the Keep Archives setting to 1 year.
 - c. Click Save.
5. Optional: Click +Management Class to add a management class.
 - a. Make selections for the basic settings, and click Add.
 - b. Customize more settings in the new management class. For backup services, make selections in the following columns: Backup Destination, Backups, Keep Extra Backups, Deleted Backups, and Keep Deleted Backups. For archive services, make selections in the Archive Destination and Keep Archives columns.
 - c. Click Save.
6. In the Default column, ensure that an appropriate management class is selected as the default. The retention settings in the default management class are applied when client operations do not specify a management class. A management class

can be specified when a client operation is run. A management class can also be specified in a client option file that is on the client system, or in a client option set that is defined on the server.

7. Activate the policy set by clicking Activate.
8. Assign client nodes to the new policy domain by either updating existing client nodes or registering new nodes.
 - o To add new clients to the policy domain, click +Client.
 - o To move an existing client to the policy domain, select the client, click Details, and then click the Properties tab. Select the new policy domain and click Save.

Data retention for the client that you assign to the policy domain is now controlled by that policy.

Requirement: If a client is running when you assign it to a new domain, you must stop and restart the client for the change to take effect.

Related tasks:


Controlling client operations through client option sets

Creating a policy by copying an existing policy

You can create new policies by copying an existing policy and then updating the parts that you want to change.

Procedure

You can create a policy by copying a policy domain, updating the management classes, and then assigning clients to the new domain.

1. On the Overview page of the Operations Center, hover over the settings icon  and click Command Builder.
2. Copy a policy domain by using the COPY DOMAIN command. For example, copy the default policy domain, STANDARD, to a new policy domain, NEWDOMAIN:

```
copy domain standard newdomain
```

This operation copies the policy domain and all associated policy sets and management classes. In this example, the operation copies the following items into the NEWDOMAIN policy domain:

- o A policy set, named STANDARD.
 - o The management class that is named STANDARD, which is in the STANDARD policy set.
 - o The copy groups that the STANDARD management class contains:
 - The backup copy group, named STANDARD
 - The archive copy group, named STANDARD
3. On the Overview page of the Operations Center, click the Services menu.
 4. Select the new policy domain and click Details. Click Policy Sets.
 5. Click the Configure toggle so that you can update the settings.
 6. Customize the settings in the management classes.
 - a. Make selections for backup services. For example, update the following items so that inactive backup versions for the clients are retained for 30 days:
 - Backups: No limit
 - Keep Extra Backups: 30 days
 - Deleted Backups: 1
 - Keep Deleted Backups: No limit
 - b. Optional: Make selections for archive services. For example, change the Keep Archives setting to 1 year.
 - c. Click Save.
 7. Optional: Make other updates and additions, such as adding a management class.
 - a. Click +Management Class to add a management class. Make selections for the basic settings, and click Add.
 - b. Customize more settings in the new management class. For backup services, make selections in the following columns: Backup Destination, Backups, Keep Extra Backups, Deleted Backups, and Keep Deleted Backups. For archive services, make selections in the Archive Destination and Keep Archives columns.
 - c. Click Save.
 8. Select the default management class that clients use, which is indicated in the Default column. Click Save. The retention settings in the default management class are applied when client operations do not specify a management class. A management class can be specified when a client operation is run. A management class can also be specified in a client option file that is on the client system, or in a client option set that is defined on the server.
 9. Activate the policy set by clicking Activate.
 10. Assign client nodes to the new policy domain by either updating existing client nodes or registering new nodes.
 - o To add new clients to the policy domain, click +Client.

- To move an existing client to the policy domain, select the client, click Details, and then click the Properties tab. Select the new policy domain and click Save.
- Data retention for the client that you assign to the policy domain is now controlled by that policy. For example, if you implemented the example in step 6, inactive backup versions for the clients are retained for 30 days by default.
- Requirement: If a client is running when you assign it to a new domain, you must stop and restart the client for the change to take effect.

Related tasks:

Controlling client operations through client option sets

Creating a policy domain

You might want to create a new policy domain for each type of client that is protected by the server. You might also want to divide responsibilities for clients among several administrators by giving them authority for specific policy domains.


About this task

Creating a new policy domain can be useful in the following circumstances:

- Applications, systems, or virtual machines require different data-retention settings. You can create a policy domain for each type of client, with a default policy that is appropriate for that type.
- Administrators are responsible for different groups of clients. For each administrator, you can create a policy domain to which you assign the clients to be managed by that administrator.

Procedure

The following steps summarize how to create a policy domain.

1. On the Overview page of the Operations Center, hover over the settings icon  and click Command Builder.
2. Define a policy domain by using the DEFINE DOMAIN command.
3. Define a policy set for the domain by using the DEFINE POLICYSET command.
4. On the Overview page of the Operations Center, click the Services menu.
5. Select the policy domain and click Details. Click Policy Sets.
6. Click the Configure toggle so that you can update the settings.
7. Click +Management Class to add a management class. Make selections for the basic settings, and click Add.
8. Optional: Customize more settings in the new management class:
 - a. For backup services, make selections in the following columns: Backup Destination, Backups, Keep Extra Backups, Deleted Backups, and Keep Deleted Backups.
 - b. For archive services, make selections in the Archive Destination and Keep Archives columns.
 - c. Click Save.
9. Optional: Click +Management Class to add more management classes.
10. In the Default column, ensure that a default management class is selected.
11. Activate the policy set by clicking Activate.
12. Assign clients to the new policy domain. From the Operations Center menu bar, click Clients.
 - To add new clients to the policy domain, click +Client.
 - To move an existing client to the policy domain, select the client, click Details, and then click the Properties tab. Select the new policy domain and click Save.

Related reference:

DEFINE DOMAIN (Define a new policy domain)

DEFINE POLICYSET (Define a policy set)

Controlling client operations through client option sets

You can use client option sets to centrally control the processing options that clients use for operations such as backup. Client option sets can help ensure that data is consistently protected according to your requirements. A client option set can override options in a local client options file, and can add options that might not be in a local client options file.

About this task

By creating and assigning client option sets, you reduce the need to update local client option files and reduce work for you or your users.

For example, you can define a client option set to specify an include-exclude list that determines what is backed up, what is excluded from backup, and what management classes are used to manage data retention. Other client options that might be useful to centrally control in a client option set are the compression and deduplication options.

You can create client option sets for clients that have similar requirements, such as clients on the same operating system, clients that use the same software, or clients that one department uses. For example, you might create client option sets for Windows workstations, or for the payroll department. After you create the client option set, you assign the client option set to all clients of the same type.

Not all client options can be specified in a client option set on the server. To learn about the client options that you can centrally control in a client option set, see [Client options that can be set by the server](#).

Procedure

1. Define a client option set by using the DEFINE CLOPTSET command. For example, to define a client option set named PAYROLLBACKUP, issue the following command:

```
define cloptset payrollbackup description='Backup options for the payroll department'
```

2. Add client options to the client option set by using the DEFINE CLIENTOPT command. For example, you want to add include and exclude options to the client option set named PAYROLLBACKUP to accomplish the following goals:
 - o Exclude temporary Internet directory files from backup operations
 - o Include for backup all files in the C:\Data directory and its subdirectories, and assign the files to the PAYCLASS management class for data retention

Issue the following commands:

```
define clientopt payrollbackup inclexcl "exclude.dir '*:\..\Temporary Internet Files'"  
define clientopt payrollbackup inclexcl "include C:\Data\..\* payclass"
```

3. To assign a client option set to a client, complete the following steps:
 - a. On the Overview page of the Operations Center, click Clients.
 - b. Select a client and click Details.
 - c. Click Properties.
 - d. In the General area, select an option set and click Save.

Related reference:

[DEFINE CLOPTSET](#) (Define a client option set name)

[DEFINE CLIENTOPT](#) (Define an option to an option set)

Related information:

[Compression client option](#)

[Deduplication client option](#)

Configuring storage

Depending on the storage functionality that you require, choose the correct type of storage media. Optimize and control your storage pools for different types of data.

- **Storage pool types**
To help you determine which storage pool type best meets your storage requirements, you should evaluate the characteristics of each storage pool type.
- **Data deduplication options**
Use inline data deduplication to deduplicate data and write the data to a container storage pool at the same time. Use postprocess data deduplication to eliminate duplicate data from sequential access (FILE) storage pools.
- **Configuring storage devices**
Configure storage devices by attaching devices, configuring device drivers, and creating the objects that represent the devices to the server.
- **Configuring a directory-container storage pool for data storage**
You can configure directory-container storage pools to use inline data deduplication to store deduplicated data.
- **Configuring a cloud-container storage pool for data storage**
You can store deduplicated data and non-deduplicated data in a cloud-container storage pool and restore the data as

required.

- Managing space in container storage pools
After you configure IBM Spectrum Protect™ and add storage, manage your data and storage pool space effectively to ensure that it operates correctly. Use container storage pools to maximize your storage space and server performance.
- Auditing a storage pool
You can schedule audit operations to identify corrupted files in storage pools.
- Auditing a storage pool container
Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.
- Storage system requirements and reducing the risk of data corruption
You can use many types of storage for the IBM Spectrum Protect server. If you use block disk storage, solid-state drives (SSD), or network-attached file systems for server storage, ensure that the storage meets requirements.

Storage pool types

To help you determine which storage pool type best meets your storage requirements, you should evaluate the characteristics of each storage pool type.

Use the following table to evaluate each type of storage pool.

Storage pool type	Description	Uses
Directory-container storage pool	A primary storage pool that a server uses to store data. Data that is stored in directory-container storage pools uses either inline data deduplication or client-side data deduplication. You can use cloud tiering to move data from directory-container storage pools to cloud-container storage pools.	Use when you want to deduplicate data inline. By using directory-container storage pools, you remove the need for volume reclamation, which improves server performance and reduces the cost of storage hardware. You cannot use this type of storage pool for storage pool backup, migration, reclamation, import or export operations.
Cloud-container storage pool	A primary storage pool that a server uses to store data. Use cloud-container storage pools to store data to an object-store based cloud storage provider. Data that is stored in cloud-container storage pools uses either inline data deduplication or client-side data deduplication.	By storing data in cloud-container storage pools, you can exploit the cost per unit advantages that clouds offer along with the scaling capabilities that cloud storage provides. You cannot use this type of storage pool for storage pool backup, migration, reclamation, import or export operations.
Random-access storage pool	A set of volumes that the server uses to store backup versions of files, files that are archive copies, and files that are migrated. Files are stored on DISK devices.	Use this type of storage pool to keep a copy of your data on DISK devices. You can migrate data into this storage pool or out of this storage pool from the following types of storage pools: <ul style="list-style-type: none"> • Random-access storage pools • Sequential-access storage pools
Sequential-access storage pool	A set of volumes that the server uses to store backup versions of files, files that are archive copies, and files that are migrated from client nodes. Files are stored on tape or FILE devices. Data that is stored in sequential-access storage pools uses both postprocess and client-side data deduplication. Restriction: Postprocess data deduplication is available in Version 7.1.2 and earlier only.	Use this type of storage pool to keep a copy of your data on FILE and TAPE devices. You can migrate data into this type of storage pool.

Storage pool type	Description	Uses
Copy storage pool	A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files).	Use copy storage pools to have a copy of active and inactive data that you can restore to a primary storage pool after a disaster or outage. You cannot use inline data deduplication, compression, replication, or data deduplication with this type of storage pool.
Container-copy storage pool	A set of tape volumes that contain a copy of deduplicated extents that reside in a directory-container storage pool. Container-copy storage pools are used only to protect the data that is stored in directory-container storage pools. Container-copy storage pools are used to repair damage in a directory-container storage pool or to restore a directory-container storage pool if a disaster occurs. Container-copy storage pools are stored on sequential media.	Use container-copy storage pools to store copies of directory-container storage pools onsite or offsite. Damaged data in directory-container storage pools can be repaired by using the deduplicated extents in a container-copy storage pool.
Active-data storage pool	A named set of storage pool volumes that contain only active versions of client backup data.	Use active-data storage pools to restore active only data to primary storage pools after a disaster or outage. By restoring active only data you can restore client data quicker and you use less bandwidth. You cannot use inline data deduplication, compression, replication, or data deduplication with this type of storage pool.

Use the following table to compare storage pool capabilities and choose the storage pool that most suits your business needs based on your storage requirements.

User goal	Directory-container storage pool	Cloud-container storage pool	Random-access storage pool	Sequential-access storage pool	Copy storage pool	Container-copy storage pool	Active-data storage pool
Protect storage pool data through node replication.	✔		✔	✔	✔		✔
Reduce storage needs by using inline compression.	✔	✔					
Reduce storage needs by using inline data deduplication.	✔	✔					

User goal	Directory-container storage pool	Cloud-container storage pool	Random-access storage pool	Sequential-access storage pool	Copy storage pool	Container-copy storage pool	Active-data storage pool
Reduce storage needs by using client-side data deduplication.	✓	✓		✓			
Reduce storage needs by using postprocess data deduplication.				✓			
Protect storage pool data through storage pool protection.	✓					✓	
Back up storage pool data by using copy storage pools on disk or tape.			✓	✓			
Store data in a cloud.		✓					
Use cloud tiering to move data from a directory-container storage pool to a cloud-container storage pool.	✓						

Data deduplication options

Use inline data deduplication to deduplicate data and write the data to a container storage pool at the same time. Use postprocess data deduplication to eliminate duplicate data from sequential access (FILE) storage pools.

You must use directory-container storage pools or cloud-container storage pools for inline data deduplication. By using directory-container or cloud-container storage pools, you reduce the need for offline reorganization, which improves server performance and reduces the cost of storage hardware. You do not use device classes or volumes with these types of storage pool.

By using postprocess data deduplication, the server identifies the data first and then removes the duplicate data to the storage pool. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance. When you remove the duplicate data, you can reclaim space in the storage pool.

For more information about postprocess data deduplication, see [Deduplicating data \(V7.1.1\)](#).

In client-side data deduplication, only compressed, deduplicated data is sent to the server. Processing is distributed between the server and the client during a backup process.

Use the following table to compare data deduplication options.

Type of data deduplication	Advantages	Disadvantages
Post-process Restriction: You can use postprocess data deduplication only with sequential access (FILE) storage pools.	<ul style="list-style-type: none"> After data deduplication, you can reclaim the storage pool. 	<ul style="list-style-type: none"> Longer processing times because the data must be identified first before the duplicate data is removed from the storage pool.
Inline Restriction: You can use inline data deduplication only with directory-container and cloud-container storage pools.	<ul style="list-style-type: none"> Deduplicates data as the data is written to a container storage pool. Reduces the need for offline reorganization which improves server performance. Reduced cost of storage hardware. 	<ul style="list-style-type: none"> Higher processor usage by the server.
Client-side	<ul style="list-style-type: none"> Processing is distributed between the server and the client during a backup process. 	<ul style="list-style-type: none"> Higher processor usage by the client. Longer elapsed time for client operations such as backup. Only compressed, deduplicated data is sent to the server.

- Defining a rule to generate data deduplication statistics
You can define a rule to generate data deduplication statistics on a regularly scheduled basis for specified nodes, node groups, and file spaces. You can generate statistics at the same time each day, or at specified intervals.

Related tasks:

- Configuring data deduplication (multisite disk solution)
- Configuring data deduplication (single-site disk solution)
- Comparing storage pools




Configuring storage devices

Configure storage devices by attaching devices, configuring device drivers, and creating the objects that represent the devices to the server.

About this task

If you are setting up a new single-site disk, multisite disk, or tape solution, you can find information about configuring storage devices in IBM Spectrum Protect™ data protection solutions.

If you are not setting up a new solution, configure and manage storage devices by following the instructions in the V7.1.1 documentation:

-   Configuring and managing storage devices
-  Configuring and managing storage devices

Configuring a directory-container storage pool for data storage

You can configure directory-container storage pools to use inline data deduplication to store deduplicated data.

Procedure

To store data in a directory-container storage pool, complete the following steps:

- Create a directory-container storage pool by completing the following steps:
 - On the Operations Center menu bar, click Storage > Storage Pools.
 - On the Storage Pools page, click + Storage Pool.
 - Complete the steps in the Add Storage Pool wizard. Select Directory for the type of container-based storage.
- After the wizard creates the storage pool, update your management classes and policy sets to use the new pool. To update a management class to use the new pool, complete the following steps:

- a. On the Operations Center menu bar, click Services.
 - b. On the Policies page, select a policy domain and click Details.
 - c. On the Details page, click the Policy Sets tab.
 - d. Click the Configure toggle. The policy sets are editable.
 - e. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.
 - f. Update one or more management classes to use the new pool by editing the Backup Destination field of the table.
 - g. Click Save.
3. Activate the changed policy set by completing the following steps:
 - a. Click Activate. Because changing the active policy set might result in data loss, a summary of the differences between the active policy set and the new policy set is displayed.
 - b. Look at the differences between corresponding management classes in the two policy sets, and consider the consequences on client files. Client files that are bound to management classes in the currently active policy set are, after activation, bound to the management classes with the same names in the new policy set.
 - c. Identify management classes in the currently active policy set that do not have counterparts in the new policy set, and consider the consequences on client files. Client files that are bound to these management classes are, after activation, managed by the default management class in the new policy set.
 - d. If the changes implemented by the policy set are acceptable, select the I understand that these updates can cause data loss check box and click Activate.
 4. Click the Configure toggle. The policy sets are no longer editable.

What to do next

To protect a directory-container storage pool, issue the PROTECT STGPOOL command. For instructions, see PROTECT STGPOOL (Protect data that belongs to a storage pool) and Copying directory-container storage pools to tape.

Linux If you are protecting a directory-container storage pool by copying the data to a remote server, and you experience network issues, see Determining whether Aspera FASP technology can optimize data transfer in your system environment.

- Copying directory-container storage pools to tape
You can protect data in a directory-container storage pool by copying the data to container-copy storage pools, which are represented by tape volumes. The tape copy is used to repair damage to a directory-container storage pool.
- Rotating tape volumes offsite when DRM is not configured
If your storage solution includes container-copy storage pools that are represented by tape volumes, but you did not configure the disaster recovery manager (DRM) function, you can follow a manual procedure to rotate the tape volumes offsite. By maintaining copies of data in offsite tape volumes, you can restore the data if a disaster occurs onsite.
- Changing the volume reclamation threshold for container-copy storage pools
By default, tape volume reclamation is enabled for container-copy storage pools. To ensure that tape volumes are used efficiently, you can change the threshold for volume reclamation.
- Reclaiming tape volumes in container-copy storage pools
You can reclaim tape volumes in container-copy storage pools without running a protection operation when you don't have time to allow both protection and reclamation operations.
- Determining whether to use container-copy storage pools for disaster protection
Determine whether container-copy storage pools meet your requirements for disaster protection.

Copying directory-container storage pools to tape

You can protect data in a directory-container storage pool by copying the data to container-copy storage pools, which are represented by tape volumes. The tape copy is used to repair damage to a directory-container storage pool.

Before you begin

Define at least one tape library to the server by using the DEFINE LIBRARY command. Provision enough tape drives and scratch volumes to meet your storage requirements. For more information about managing backup media and configuring disaster recovery manager (DRM), see Disaster recovery manager (V7.1.1).

About this task

To copy the data in directory-container storage pools to tape, the Operations Center creates a schedule to run the PROTECT STGPOOL command. When the protection schedule runs, one tape copy is created. At least one volume must be available when the protection schedule runs. Otherwise, the operation fails.

You can create up to two tape copies, but you must use the command-line interface to create a second container-copy storage pool. One tape copy can be taken to an offsite disaster recovery location. The other copy can be kept onsite to expedite recovery from less-critical failures.

Restrictions:

- Virtual tape libraries are not supported, regardless of which library type is defined. Only physical tape is supported.
- Container-copy storage pools can be used to repair minor to moderate storage pool damage, which includes damaged containers or directories. Container-copy storage pools can also be used for disaster protection, but you must ensure that recovery times meet your requirements. For more information, see [Determining whether to use container-copy storage pools for disaster protection](#).
- You cannot use replication to target a container-copy storage pool.
Tip: You can create a tape copy of the directory-container storage pool data at a disaster recovery site by using this procedure to create a container-copy storage pool on the target replication server. Then, schedule the PROTECT STGPOOL and REPLICATE NODE commands to run on the source replication server to protect your data to the target replication server.
- You cannot use the following procedure if the directory-container storage pool already has an associated container-copy storage pool. To create a second container-copy storage pool, follow the instructions in step 5.

If you created a container-copy storage pool as part of the Add Storage Pool wizard, you do not have to use this procedure. When you completed the wizard, the Operations Center configured the container-copy storage pool and a protection schedule.

Procedure

To configure storage pool protection to tape for an existing directory-container storage pool, complete the following steps:

1. On the Operations Center menu bar, click Storage > Storage Pools.
2. On the Storage Pools page, select the directory-container storage pool that you want to protect to tape.
3. Click More > Add Container-Copy Pool.
4. Follow the instructions in the Add Container-Copy Pool window to schedule protection to tape.
5. After you complete the previous steps, you can add a second container-copy storage pool by using the command-line interface. Optionally, complete the following steps to add a container-copy storage pool:
 - a. Create a container-copy storage pool by issuing the DEFINE STGPOOL command.
 - b. Assign the container-copy storage pool to the directory-container storage pool by issuing the UPDATE STGPOOL command for the directory-container pool.

Results

After you complete the configuration, data in the directory-container storage pool is copied to a container-copy storage pool based on the defined protection schedule.

What to do next

1. If you created a tape copy to store offsite, enable the offsite container-copy storage pool for DRM operations by issuing the SET DRMCOPYCONTAINERSTGPOOL command. Ensure that the tape volumes are added to your offsite tape rotation schedules. If DRM is not configured, you must do so or use the alternative method to rotate tapes offsite. For instructions about the alternative method, see [Rotating tape volumes offsite when DRM is not configured](#). To verify that offsite container-copy storage pools are enabled for DRM, use the QUERY DRMSTATUS command.

For instructions about configuring DRM, see [Disaster recovery manager \(V7.1.1\)](#).

2. Confirm that the reclamation threshold for your container-copy storage pool meets your requirements.

By default, tape volume reclamation is enabled for new container-copy storage pools that are created by using the Operations Center. Volume reclamation occurs when the reclamation threshold for the container-copy storage pool is less than 100%. However, tape volumes are not a candidate for reclamation until they are 75% full. Be careful when you use reclamation with container-copy storage pools that have offsite volumes. When an offsite volume becomes eligible for reclamation, in effect the server moves the extents on the volume back to the onsite location. If a disaster occurs onsite, the server can obtain extents from the offsite volume if the restored database refers to extents on the offsite volume. To prevent volumes from being rewritten immediately after all extents are deleted, use the REUSEDELAY parameter to specify a value that is greater than 0. The Operations Center sets the reclamation threshold at 60% for onsite container-copy storage pools.

For instructions about changing the reclamation threshold, see [Changing the volume reclamation threshold for container-copy storage pools](#).

3. Protect the metadata for your container-copy storage pool.

When the protection schedule runs, data extents in the container-copy storage pools are copied to tape volumes without the associated metadata. This metadata is required to restore the tape copies. To protect the metadata, you must separately back up the server database along with its volume history, server options, and device configuration files. If you use reclamation with container-copy storage pools that have offsite tape volumes, ensure that the following requirements are met to provide disaster recovery protection:

- o Database backup operations run after storage pool protection schedules and DRM move schedules finish.
- o All database backup volumes and DRM volumes are taken offsite together.

For instructions about backing up the server database and related files, see [Defining schedules for server maintenance activities](#).

4. Optionally, change the protection schedule for a directory-container storage pool that has one or more associated container-copy storage pools by using the UPDATE SCHEDULE command. The schedule that is created by the Operations Center is named CONTAINER_COPY.

Related concepts:

Data storage in container-copy storage pools

Related tasks:

Determining whether to use container-copy storage pools for disaster protection

Related reference:

DEFINE LIBRARY (Define a library)

PROTECT STGPOOL (Protect data that belongs to a storage pool)

UPDATE SCHEDULE (Update an administrative schedule)

QUERY DRMSTATUS (Query disaster recovery manager system parameters)

Rotating tape volumes offsite when DRM is not configured

If your storage solution includes container-copy storage pools that are represented by tape volumes, but you did not configure the disaster recovery manager (DRM) function, you can follow a manual procedure to rotate the tape volumes offsite. By maintaining copies of data in offsite tape volumes, you can restore the data if a disaster occurs onsite.

Procedure

1. Check out the storage volume that must be rotated offsite by using the CHECKOUT LIBVOLUME command.
2. Update the volume to indicate that it is moved offsite by using the UPDATE VOLUME command and specifying ACCESS=OFFSITE. Optionally, indicate the offsite location by using the LOCATION parameter. For example, specify LOCATION=SITE1.
3. Reclaim space by taking one of the following actions:
 - o To reclaim space without protecting the storage pool, run the PROTECT STGPOOL command and specify TYPE=LOCAL and RECLAIM=ONLY.
 - o To reclaim space while protecting the storage pool, run the PROTECT STGPOOL command without specifying the RECLAIM parameter.
4. Monitor the volume by using the QUERY VOLUME command. If the volume is shown to be unavailable and empty, return the volume onsite and check it into the library by using the CHECKIN LIBVOLUME command.
5. Update the volume by using the UPDATE VOLUME command and specifying ACCESS=READWRITE.

Related reference:

CHECKOUT LIBVOLUME (Check a storage volume out of a library)

PROTECT STGPOOL (Protect data that belongs to a storage pool)

UPDATE VOLUME (Change a storage pool volume)

Changing the volume reclamation threshold for container-copy storage pools

By default, tape volume reclamation is enabled for container-copy storage pools. To ensure that tape volumes are used efficiently, you can change the threshold for volume reclamation.

Procedure

1. On the Operations Center Overview page, click Storage > Storage Pools.
2. Select the storage pool and click Details, and then Properties.
3. In the Reclamation section, set the reclamation percentage and click Save.

Tip: Alternatively, change the reclamation threshold by issuing the UPDATE STGPOOL command and specifying the RECLAIM parameter. For details about the RECLAIM parameter, see the commands for defining and updating container-copy storage pools.

Restriction: You cannot use the RECLAIM STGPOOL command to reclaim volumes in container-copy storage pools. For details about reclaiming volumes in container-copy storage pools, see the RECLAIM parameter in the PROTECT STGPOOL command.

Reclaiming tape volumes in container-copy storage pools

You can reclaim tape volumes in container-copy storage pools without running a protection operation when you don't have time to allow both protection and reclamation operations.

About this task

When you issue the PROTECT STGPOOL command and the target storage pool is a container-copy storage pool, both protection and reclamation operations run by default. The preferred practice is to allow both protection and reclamation operations to run. However, to save time, you can run only the storage pool protection operation or only reclamation, or you can limit the number of tape volumes that are reclaimed. Use this procedure only when you have to reclaim tape volumes quickly or when you have to reclaim a limited number of tape volumes.

Procedure

To reclaim tape volumes without running a storage pool protection operation, complete the following steps:

1. Optional: To maximize the amount of space that is reclaimed, start the inventory expiration process by issuing the EXPIRE INVENTORY command.
2. Determine whether you want reclamation to run to completion or limit the number of tape volumes that are reclaimed.
3. To run reclamation to completion, issue the PROTECT STGPOOL command and specify the TYPE=LOCAL and RECLAIM=ONLY parameters. For example, to reclaim space in a local container-copy storage pool that is defined as the target protection pool for SPOOL1, issue the following command:

```
protect stgpool spool1 type=local reclaim=only
```

4. To reclaim a limited number of tape volumes, complete the following steps:
 - a. Set a reclamation limit for the container-copy storage pool by issuing the UPDATE STGPOOL command and specifying the RECLAIMLIMIT parameter. This parameter limits the number of volumes in the container-copy storage pool that are reclaimed.
 - b. Issue the PROTECT STGPOOL command and specify the TYPE=LOCAL parameter along with either the RECLAIM=YESLIMITED or RECLAIM=ONLYLIMITED parameter.

Tip: When you specify RECLAIM=YESLIMITED, both reclamation and storage pool protection operations run when the PROTECT STGPOOL command is issued. When you specify RECLAIM=ONLYLIMITED, reclamation is the only operation that runs. When you specify either of these values, reclamation runs only until it reaches the reclamation limit that is defined for the container-copy storage pool. The reclamation limit is defined with the RECLAIMLIMIT parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.

For example, to reclaim a limit of five tape volumes in a container-copy storage pool that is named CCPool1 without running a protection operation on the source directory-container storage pool that is named SPOOL1, issue the following commands:

```
update stgpool ccpool1 reclaimlimit=5  
protect stgpool spool1 type=local reclaim=onlylimited
```

For example, to protect a storage pool that is named SPOOL1 and to reclaim a maximum of 10 tape volumes in the associated container-copy storage pool, issue the following commands:

```
update stgpool spool1 reclaimlimit=10  
protect stgpool spool1 type=local reclaim=yeslimited
```

Results

Reclamation processing for the container-copy storage pool is completed. The storage pool protection operation did not run, so data in the directory-container storage pool that was updated since the last protection operation is not protected.

What to do next

1. Protect the data in the directory-container storage pool to the container-copy storage pool by issuing the PROTECT STGPOOL command and specifying the TYPE=LOCAL parameter. The protection process runs with the default RECLAIM=YES parameter. The protection operation takes less time because reclamation already ran. For example, to protect the data in a directory-container storage pool that is named SPOOL1, issue the following command:

```
protect stgpool spool1 type=local
```

Alternatively, protect the data in a directory-container storage pool that is named SPOOL1 without running reclamation by issuing the following command:

```
protect stgpool spool1 type=local reclaim=no
```

2. Back up the server database and run scheduled maintenance operations. For instructions, see Defining schedules for server maintenance activities.

Related reference:

PROTECT STGPOOL (Protect data that belongs to a storage pool)

DEFINE STGPOOL (Define a container-copy storage pool)

UPDATE STGPOOL (Define a container-copy storage pool)

EXPIRE INVENTORY (Manually start inventory expiration processing)

Determining whether to use container-copy storage pools for disaster protection

Determine whether container-copy storage pools meet your requirements for disaster protection.

About this task

You can create an offsite copy of your container-copy storage pool for disaster recovery protection or to satisfy regulatory and business requirements for offsite tape copies. Before you decide to use offsite tape copies for disaster protection, carefully consider whether the solution meets your recovery time objective.

Using container-copy storage pools for disaster recovery is suitable when the amount of data in your environment is equal to or less than the following values:

- 200 TB of total managed data
- 50 TB of back-end data
- 37 TB of front-end data

Total managed data

All data that is stored in the directory-container storage pool on the server. This includes active and inactive versions of the data. The number of versions is determined by retention policies.

Back-end data

All data that is stored in the container-copy storage pool.

Front-end data

The current active data that is stored in the container-copy storage pool. This is the active data that is used to restore data on client nodes. In a disaster, all or part of the front-end data is required to reestablish production. Front-end data is a percentage of total managed data and is less than or equal to the total managed data, depending on the policy settings in use.

To recover from a disaster within 48 hours, the system environment at the recovery site must meet the minimum hardware requirements for the actions in the following table.

Action	Time required	Minimum requirements
--------	---------------	----------------------

Action	Time required	Minimum requirements
Configure a new IBM Spectrum Protect™ server at a disaster recovery site. To configure the new server, you must complete the following steps: <ol style="list-style-type: none"> 1. Provision disks for the server. 2. Restore the server from backup. 3. Start the server. 4. Update the storage and device configurations. 	Time to restore the server: 6 hours	Use a solid-state drive (SSD) for the server database, with the following requirements: <ul style="list-style-type: none"> • A minimum of 100 MB per second average combined read/write throughput • A minimum of 12,862 average input/output operations per second (IOPS)
Audit the directory-container storage pool and repair the data from tape. Tip: If the system meets the minimum hardware requirements, you can repair up to 50 TB of back-end data within 48 hours.	Time to audit the storage pool: 2 hours Time to repair the storage pool by using a tape copy: 28 hours Note: The time estimate applies if you have a maximum of 200 TB of total managed data in the storage pool.	Use Nearline SAS (NL-SAS) drives, as in a medium blueprint server configuration, with a minimum of 700 MB per second write performance to storage pool disk. Use new generation tape technology such as LTO-7 or better, with a minimum of six drives to allow concurrent read operations from tape volumes.
Restore data on client nodes. Tip: If the system meets the minimum hardware requirements, you can restore up to 37 TB of front-end data within 48 hours.	Time for client restore operations: 12 hours	Use NL-SAS drives, as in a medium blueprint server configuration, with a minimum of 10 restore sessions achieving 3102 GB per hour.

Procedure

1. Estimate the disaster recovery time for your environment by using the following table. Determine whether the recovery time meets your requirements.

Table 1. Recovery time estimate for differing amounts of total managed data

Recovery time objective	Total managed data (TB)	Number of hours to repair a directory-container storage pool (First Byte Restored)	Hours until client nodes are restored (Disaster Recovery complete)
Up to 1 day	25	10	12
	50	13	16
	75	17	22
Up to 2 days	100	20	26
	200	34	46
Up to 4 days	300	48	66
	400	62	86
More than 4 days	500	76	106

Notes:

- o Achievable rates are highly dependent on the workload and the configured environment.
 - o The front-end data percentage is relative to the total managed data. Increasing the amount of front-end data increases the total recovery time. Decreasing the amount of front-end data decreases the total recovery time.
2. Estimate the recovery time for your environment by using the following formulas:

- o Estimate the value **Hours until directory-container storage pool is repaired (First Byte Restored)**:

$$\text{Time to Client First Byte Restore} = 6 \text{ hours} + 14 \text{ hours for every } 100 \text{ TB of Total Managed Data}$$

- o Estimate the value **Hours until client nodes are restored (Disaster Recovery complete)**:

Time to Client Restore Complete =
Time to Client First Byte Restore + ((Total Managed Data * Front-End Data) / Restore Rate)

Restore Rate: The rate at which clients can restore data from the server back to their local computer or storage device.

3. Complete test procedures for disaster recovery to ensure that container-copy storage pools can be used to restore your environment in a time frame that meets your requirements.

Related reference:

Repairing storage pools after a disaster

Configuring a cloud-container storage pool for data storage

You can store deduplicated data and non-deduplicated data in a cloud-container storage pool and restore the data as required.

Before you begin

Review the requirements and restrictions that apply to cloud-container storage pools.

You can configure cloud-container storage pools to use one of the following service providers and protocols:

- Amazon Web Services (AWS) with Simple Storage Service (S3)
- Microsoft Azure
- IBM® Cloud Object Storage with S3
- IBM Cloud Object Storage with S3 and IBM Cloud
- IBM Cloud Object Storage with Swift and IBM Cloud
- OpenStack with Swift using Keystone Version 1 or 2

Restriction: Cloud-container storage pools are not supported on the Linux on System z operating system.

Obtain configuration information and specify a device class by completing the following steps:

1. Obtain the configuration information for your cloud service provider:
 - Amazon with S3 (off-premises)
 - Microsoft Azure
 - IBM Cloud Object Storage with S3 (off-premises, with IBM Cloud)
 - IBM Cloud Object Storage with Swift (off-premises, with IBM Cloud)
 - IBM Cloud Object Storage with S3 (on-premises)
 - OpenStack with Swift (on-premises or off-premises)
2. Specify a device class to use for database backup operations. When you use encryption for cloud-container storage pools, the server master encryption key is used to protect the cloud encryption key in a database backup.
 - a. On the Operations Center menu bar, click Servers.
 - b. Select a server row and click Back Up.
 - c. Select a device class to use for database backup operations and click Back Up.

Tip: Alternatively, use the SET DBRECOVERY command to specify a device class for the database backup.

Procedure

To store data in a cloud-container storage pool, complete the following steps:

1. Create a cloud-container storage pool. You must provide configuration information that identifies the cloud service.
 - a. On the Operations Center menu bar, click Storage > Storage Pools.
 - b. On the Storage Pools page, click + Storage Pool.
 - c. Complete the steps in the Add Storage Pool wizard. Select On-premises cloud or Off-premises cloud for the type of container-based storage.
2. Update your management classes and policy sets to use the new storage pool. To update a management class to use the new storage pool, complete the following steps:
 - a. On the Operations Center menu bar, click Services.
 - b. On the Policies page, select a policy domain and click Details.
 - c. On the Details page, click the Policy Sets tab.
 - d. Click the Configure toggle. The policy sets are editable.
 - e. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.

- f. Update one or more management classes to use the new storage pool by editing the Backup Destination field of the table.
 - g. Click Save.
3. Activate the changed policy set by completing the following steps:
 - a. Click Activate. Because changing the active policy set might result in data loss, a summary of the differences between the active policy set and the new policy set is displayed.
 - b. Look at the differences between corresponding management classes in the two policy sets, and consider the consequences on client files. Client files that are bound to management classes in the currently active policy set are, after activation, bound to the management classes with the same names in the new policy set.
 - c. Identify management classes in the currently active policy set that do not have counterparts in the new policy set, and consider the consequences on client files. Client files that are bound to these management classes are, after activation, managed by the default management class in the new policy set.
 - d. If the changes implemented by the policy set are acceptable, select the I understand that these updates can cause data loss check box and click Activate.
 4. Click the Configure toggle. The policy sets are no longer editable.
 5. To take advantage of local storage, create a storage pool directory for this storage pool using the DEFINE STGPOOLDIRECTORY command. For more information, see Optimizing performance for cloud object storage.

Related tasks:

- Preparing to configure cloud-container storage pools for AWS with S3 (off premises)
- Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with S3 (on premises)
- Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with S3 (off premises)
- Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with Swift (off premises)
- Preparing to configure cloud-container storage pools for OpenStack with Swift
- Encrypting data for cloud-container storage pools
- Optimizing performance for cloud object storage

Related reference:

- SET DBRECOVERY (Set the device class for automatic backups)

Preparing to configure cloud-container storage pools for AWS with S3 (off premises)

Before you configure cloud-container storage pools to use Amazon Web Services (AWS) off premises with the Simple Storage Service (S3) protocol, you must obtain information from Amazon that is required for the configuration process.

About this task

AWS account credentials are different from Amazon account credentials. Use the credentials for your AWS account when you configure storage pools in the Operations Center or with the DEFINE STGPOOL command.

AWS uses *buckets* to store data. AWS buckets are used in the same manner as containers in a cloud-container storage pool. IBM Spectrum Protect™ automatically creates a bucket in Amazon for an instance of IBM Spectrum Protect, and that bucket is shared by all pools for that instance.

Restriction: The following restrictions apply.

- Edit an AWS bucket only with IBM Spectrum Protect, and do not change the data in the bucket or edit the configuration settings for the bucket.
- For off-premises cloud-container storage pools that use AWS with the Amazon S3 protocol, data is encrypted by default. However, the IBM Spectrum Protect server does not support encryption of the data by using AWS bucket policies.

Procedure

1. Sign up for an AWS account by going to the Amazon S3 page and clicking Create an AWS Account.
2. Obtain your AWS credentials:
 - a. Go to the Amazon S3 page and click Sign In to the Console.
 - b. Select your name and select Security Credentials.
 - c. Go to the Access Keys section to locate the Access Key ID and the Secret Access Key fields. Record the values so that you can use them when you configure storage pools.
3. If you plan to configure storage pools by using the Add Storage Pool wizard in the Operations Center, use the following values for the parameters:

- o Cloud type: `Amazon - S3 API`
 - o Access key ID: `access_key_id`
 - o Secret access key: `secret_access_key`
 - o Region: Select the region endpoint that best fits your location, based on the AWS Regions and Endpoints page. If you select `Other`, specify a region endpoint URL in the URL field, and include the protocol, usually `https://`. Typically, you can use the region that is closest to your physical location for the Region parameter. Because an Amazon bucket exists in only one region, you can specify only one endpoint URL for a region. If you require a GovCloud region, specify a URL from the AWS GovCloud (US) Endpoints page.
Warning: Be sure to use only the AWS endpoint URL for the Region value, such as `https://s3-us-west-1.amazonaws.com`. Do not use the static website hosting URL for this value.
 - o Bucket name: Use the default bucket name generated by the server, or specify a new bucket name.
4. To define the cloud-container storage pool, issue the `DEFINE STGPOOL` command with the following values:
- o `CLOUDTYPE: S3`
 - o `IDENTITY: access_key_id`
 - o `PASSWORD: secret_access_key`
 - o `CLOUDURL`: Specify the region endpoint URL that best fits your location, based on the AWS Regions and Endpoints page.

Typically, you can use the region that is closest to your physical location for the `CLOUDURL` parameter. If you require a GovCloud region, specify a URL from the AWS GovCloud (US) Endpoints page.

Warning: Be sure to use only the AWS endpoint URL for the `CLOUDURL` value, such as `https://s3-us-west-1.amazonaws.com`. Do not use the static website hosting URL for this value.

What to do next

Configure cloud-container storage pools for AWS by following the instructions in [Configuring a cloud-container storage pool for data storage](#).

Configuring an Amazon S3 compatible device as a cloud-container storage pool

You can configure a storage device that is compatible with the Amazon Simple Storage Service (S3) protocol so that the device can be used as an IBM Spectrum Protect™ cloud-container storage pool.

About this task

Amazon S3 uses *buckets* to store data. You must create a bucket on the S3 compatible storage device for use by an IBM Spectrum Protect server. After you create the bucket, use the credentials from the account on your Amazon S3 compatible cloud object storage device when you configure storage pools with the `DEFINE STGPOOL` command.

Restriction: Do not change the data in the bucket or edit the configuration settings for the bucket.

Procedure

1. Create a bucket on the cloud object storage device. Follow the instructions in the device documentation.
2. Create a user account on the cloud object storage device. The account is used by IBM Spectrum Protect to access the device by using the access key ID and secret access key. Ensure that the account has permissions to store data in and delete data from the bucket that you created in Step 1. Record the access key ID and the secret access key values so that you can use them when you configure storage pools.
3. Identify the URL value that will be used by IBM Spectrum Protect to access the cloud object storage device. For instructions, see the documentation for your cloud object storage device.
4. To define the cloud-container storage pool, issue the `DEFINE STGPOOL` command with the following values:
 - o `CLOUDTYPE: S3`
 - o `IDENTITY: access_key_id`
 - o `PASSWORD: secret_access_key`
 - o `CLOUDURL`: `http://cloud_object_storage_endpoint_IP_address` or `https://cloud_object_storage_endpoint_IP_address`. If you use more than one endpoint, list the endpoint IP addresses separated by a vertical bar (`|`), with no spaces, as shown in the following example:

```
CLOUDURL=endpoint_URL1|endpoint_URL2|endpoint_URL3
```

- o BUCKETNAME: *name_of_bucket_on_device*

To optimize performance, use multiple endpoints or a load balancer.

What to do next

Configure cloud-container storage pools in a similar way as you would configure a cloud-container storage pool for IBM Cloud Object Storage by following the instructions in [Configuring a cloud-container storage pool for data storage](#).

Preparing to configure cloud-container storage pools for Microsoft Azure (off premises)

Before you configure cloud-container storage pools to use the Microsoft Azure cloud computing system, you must obtain information for the configuration process from Microsoft.

About this task

IBM Spectrum Protect™ supports the following Azure storage tiers:

- *Hot storage tier* for data that is accessed frequently
- *Cool storage tier* for data that is accessed less frequently

You can use a cool storage tier for cost-effective, long-term storage. However, it is more costly to restore data from a cool storage tier than from a hot storage tier.

Procedure

1. Sign up for a Microsoft Azure account by going to the Azure portal and creating an account.
2. Create a storage account. Typically, select the location that is nearest to your IBM Spectrum Protect server for the storage account location.
3. Obtain your Azure credentials:
 - a. Go to the Azure portal and click Storage accounts.
 - b. Open the new storage account, go to the container section of the Blob Service pane, and record the blob service endpoint value so that you can use it when you configure storage pools. The blob service endpoint looks like these examples: `https://name.blob.core.windows.net` and `http://name.blob.core.windows.net`.
 - c. Create a shared access signature (SAS) token by opening the Shared access signature tab and completing the fields. Ensure that the Allowed services section includes Blob and that the Allowed resource types section includes Container and Object. Ensure that the SAS token has read, write, delete, list, add, and create permissions. Click Generate SAS.
 - d. Record the SAS token value so that you can use it when you configure storage pools. IBM Spectrum Protect does not monitor the SAS token expiration date, so ensure that you select a date that best suits your needs. If the token expires, the IBM Spectrum Protect server loses access to the storage account until you provide a new SAS token. Tip: If you would like to less frequently update the SAS token, set an expiration date that is several years away. Also, ensure that you verify the start date and time fields.
4. If you plan to configure storage pools by using the Add Storage Pool wizard in the Operations Center, use the following values for the parameters:
 - o Cloud type: `Azure`
 - o SAS token: `SAS_token_value`. Look for a string that is similar to this example:


```
?sv=2016-05-31&ss=b&srt=sco&sp=rwdlac&se=2017-04-05T18:26:12Z&st=2017-04-05T10:26:12Z&spr=https&sig=XUangS%2FcXXXXXXXXXXXXXXXXXXXXXXXXXXXXXELsuWp106Cmq7o%3D
```
 - o Blob service endpoint: Specify the blob service endpoint from your Azure storage account, for example, `https://name.blob.core.windows.net` or `http://name.blob.core.windows.net`.
5. If you plan to configure storage pools by using the DEFINE STGPOOL command, use the following values for the command parameters:
 - o CLOUDTYPE: `Azure`
 - o PASSWORD: `SAS_token_value`. Look for a string that is similar to this example:


```
?sv=2016-05-31&ss=b&srt=sco&sp=rwdlac&se=2017-04-05T18:26:12Z&st=2017-04-05T10:26:12Z&spr=https&sig=
```

- o CLOUDURL: Specify the blob service endpoint from your Azure storage account, for example, `https://name.blob.core.windows.net` or `http://name.blob.core.windows.net`.

What to do next

Configure cloud-container storage pools for Azure by following the instructions in [Configuring a cloud-container storage pool for data storage](#).

Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with Swift (off premises)

Before you configure cloud-container storage pools to use IBM® Cloud Object Storage and IBM Cloud off premises using Swift, you must obtain configuration information from the [IBM Cloud Object Storage page](#).

About this task

Use the credentials from your IBM Cloud account when you configure the storage pools in the Operations Center or with the `DEFINE STGPOOL` command.

Procedure

1. Create an IBM Cloud account by following the instructions in the [IBM Cloud Docs](#) section.
2. Obtain your IBM Cloud credentials:
 - a. Go to the [IBM Cloud Object Storage page](#) and log in with your account credentials.
 - b. Select the account and cluster that you want to configure.
 - c. In the Account section, click [View Credentials](#)
 - d. In the Account Credentials section, locate the Public Authentication Endpoint, Username, and API Key fields. Record the values in those fields so that you can use them when you configure storage pools.
3. If you plan to configure storage pools by using the Add Storage Pool wizard in the Operations Center, use the following values for the parameters:
 - o Cloud type: `IBM Cloud Object Storage - Swift API`
 - o User name: `username`
 - o Password: `API_key`
 - o URL: `public_authentication_endpoint`
4. If you plan to configure storage pools by using the `DEFINE STGPOOL` command, use the following values for the command parameters:
 - o `CLOUDTYPE`: `IBMCLOUDSWIFT`
 - o `IDENTITY`: `username`
 - o `PASSWORD`: `API_key`
 - o `CLOUDURL`: `public_authentication_endpoint`

What to do next

Configure cloud-container storage pools for IBM Cloud by following the instructions in [Configuring a cloud-container storage pool for data storage](#).

Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with S3 (off premises)

You can set up cloud-container storage pools to use IBM® Cloud Object Storage off premises with the Simple Storage Service (S3) protocol.

About this task

The off-premises implementation of IBM Cloud Object Storage is managed through IBM Cloud. In this setup, only the owner of the IBM Cloud account can create buckets and administrators.

Use the credentials from your IBM Cloud account when you configure the storage pools in the Operations Center or with the DEFINE STGPOOL command. For more information, see the IBM Cloud Storage page. To use this configuration, select Cloud Object Storage - S3 API from the IBM Cloud Order Object Storage page.

Procedure

1. Log into the IBM Cloud Customer Portal.
2. Click the Storage menu and select Object Storage.
3. From the Object Storage page, select an S3 account.
4. From the Cloud Object Storage page, click Manage Buckets and then click the + symbol to create the bucket you want to use with the new cloud-container storage pool.
5. Click Show Credentials to create administrator credentials for your new bucket.
6. Click Add Credential.
7. Locate the Access Key ID, the Secret Access Key, and the Public Authentication Endpoint. Record the values in those fields so that you can use them when you configure storage pools. If you are inside the IBM Cloud network, you can use a private authentication endpoint.
8. To configure storage pools by using the Add Storage Pool wizard in the Operations Center, select Off-premises cloud. Use the following values for the parameters:
 - o Cloud type: IBM Cloud Object Storage - S3 API
 - o Access key ID: *access_key_ID*
 - o Secret access key: *secret_access_key*
 - o Bucket name: *bucket_name* (from step 4)
 - o URL: *us-geo_authentication_endpoint*

Note: Only one cloud provider endpoint is needed with this configuration. If all of your servers are inside the IBM Cloud network, you can use a private authentication endpoint.
9. If you configure storage pools by using the DEFINE STGPOOL command, use the following values for the command parameters:
 - o CLOUDTYPE: S3
 - o IDENTITY: *access_key_ID*
 - o BUCKETNAME: *bucket_name* (from step 4)
 - o PASSWORD: *secret_access_key*
 - o CLOUDURL: *us-geo_authentication_endpoint*

Note: Only one cloud provider endpoint is needed with this configuration. If all of your servers are inside the IBM Cloud network, you can use a private authentication endpoint.

What to do next

Configure cloud-container storage pools for IBM Cloud object storage by following the instructions in Configuring a cloud-container storage pool for data storage.

Preparing to configure cloud-container storage pools for IBM Cloud Object Storage with S3 (on premises)

Before you configure cloud-container storage pools to use IBM® Cloud Object Storage on premises with S3, you must set up an IBM Cloud Object Storage vault template and an IBM Cloud Object Storage user account, and then obtain configuration information.

About this task

Restriction:

To use IBM Cloud Object Storage on premises with S3, ensure that your version of IBM Cloud Object Storage is compatible with your version of IBM Spectrum Protect™.

For IBM Spectrum Protect Version 8.1.4, IBM Cloud Object Storage V3.8.3 or later is required.

IBM Cloud Object Storage vaults are used in the same manner as containers in a cloud-container storage pool. Set up a vault template to quickly create vaults with your preferred settings.

After you create a vault template, use the credentials from your IBM Cloud Object Storage user account to configure the storage pools in the Operations Center or with the DEFINE STGPOOL command. The server uses the Simple Storage Service (S3) protocol

to communicate with IBM Cloud Object Storage.

Tip: You can skip the first four steps in the procedure if you want to configure an existing vault by using the BUCKETNAME parameter in the DEFINE STGPOOL or UPDATE STGPOOL commands.

Procedure

1. Create a vault template:
 - a. Log in to IBM Cloud Object Storage and click the Configure tab.
 - b. In the dsNet navigation pane, expand Storage Pools.
 - c. Select the IBM Cloud Object Storage storage pool where you want to create the vault template, and click the Storage Pool link in the General section.
 - d. In the Vault Templates section, click Create Vault Template.
 - e. Select the settings for the default vault template. You might be able to optimize performance by not selecting the Enable SecureSlice Technology or the Name Index Enabled options, and selecting the Recovery Listing Enabled option.
 - f. In the Deployment section, select the access pool or pools that you want to use for the template and click Save.
2. Set the vault template as the default for your IBM Cloud Object Storage dsNet:
 - a. Click the Configure tab.
 - b. In the Default Vault Template Configuration section, click Configure.
 - c. Select a vault template to use as the default, and click Update to set that template as the default.
3. If this is your first time configuring a vault template, enable the vault provisioning role so you can create new vaults:
 - a. Click the Administration tab.
 - b. In the Provisioning API Configuration section, click Configure.
 - c. Select Create Only or Create and Delete to let users create new vaults using the Provisioning API.
 - d. Click Update to save the settings.
4. Use an IBM Cloud Object Storage account with administration authority to create a user account on the IBM Cloud Object Storage instance in your environment. Ensure that the new user account has the Vault Provisioner role.
5. Click the Security tab and select the new user account.
6. Generate an access key for the new user:
 - a. In the Access Key Authentication section, click Change Keys.
 - b. On the Edit Access Keys page, click Generate New Access Key.
 - c. Click Back.
7. In the Access Key Authentication section, locate the Access Key ID and Secret Access Key values. Record the values so that you can use them when you configure storage pools.
8. Locate the URL value:
 - a. Click the Configure tab.
 - b. In the dsNet navigation pane, expand the Devices and Accesser sections.
 - c. Select the IBM Cloud Object Storage accesser. Verify that the accesser belongs to an access pool to which the default vault template is deployed.
 - d. In the Device Configuration section for the accesser, record the IP Address value so that you can use it when you configure storage pools. Use `http://` before the IP address value to prevent certificate security errors.
9. If you configure storage pools by using the Add Storage Pool wizard in the Operations Center, use the following values for the parameters:
 - o Cloud type: IBM Cloud Object Storage - S3 API
 - o Access key ID: *access_key_ID*
 - o Secret access key: *secret_access_key*
 - o Bucket name: Use the default bucket name generated by the server, or specify a new bucket name.
 - o URL: `http://Cloud_Object_Store_accesser_IP_address`
Important: If you use more than one accesser, type an accesser IP address and then press Enter to add additional IP addresses. Use multiple accessers or a load balancer for optimal performance.
10. If you configure storage pools by using the DEFINE STGPOOL command, use the following values for the command parameters:
 - o CLOUDTYPE: S3
 - o IDENTITY: *access_key_ID*
 - o PASSWORD: *secret_access_key*
 - o CLOUDURL: `http://Cloud_Object_Store_accesser_IP_address`
Important: If you use more than one accesser, list the accesser IP addresses separated by a vertical bar (|), with no spaces, such as `CLOUDURL=<accesser_URL1>|<accesser_URL2>|<accesser_URL3>`. Use multiple accessers or a load balancer for optimal performance.

What to do next

Configure cloud-container storage pools for IBM Cloud Object Storage by following the instructions in Configuring a cloud-container storage pool for data storage.

Preparing to configure cloud-container storage pools for OpenStack with Swift

Before you configure cloud-container storage pools to use OpenStack on premises or off premises with Swift, you must obtain configuration information from the OpenStack Swift computer.

About this task

Restriction: You must use OpenStack Swift Release Series Juno, Service Project keystone, application programming interface (API) Version 1 or 2.

Use the credentials from your OpenStack Swift account when you configure the storage pools by using the Operations Center or the DEFINE STGPOOL command.

Procedure

1. Create an OpenStack Swift account by following the instructions in the OpenStack Swift documentation.
2. Obtain your OpenStack Swift credentials:
 - a. On the OpenStack Swift computer, type the following command:

```
swift auth -v
```
 - b. In the output, locate the `OS_AUTH_URL`, the `OS_TENANT_NAME`, the `OS_USERNAME`, and the `OS_PASSWORD` values. Record the values so that you can use them when you configure storage pools.
3. If you plan to configure storage pools by using the Add Storage Pool wizard in the Operations Center, use the following values for the parameters:
 - o Cloud type: `OpenStack Swift`
 - o User name: `OS_TENANT_NAME:OS_USERNAME`
 - o Password: `OS_PASSWORD`
 - o URL: `OS_AUTH_URL`
4. If you plan to configure storage pools by using the DEFINE STGPOOL command, use the following values for the command parameters:
 - o CLOUDTYPE: `SWIFT` or `V1SWIFT`
 - o IDENTITY: `OS_TENANT_NAME:OS_USERNAME`
 - o PASSWORD: `OS_PASSWORD`
 - o CLOUDURL: `OS_AUTH_URL`
5. If you plan to use a specific tenant or user name, record the values in the following format: `TENANT_NAME:USERNAME`.
6. To prevent data loss, configure OpenStack Swift to create replicas of the data that is written to its object storage. For more information, see the OpenStack Swift documentation.

What to do next

Configure cloud-container storage pools for OpenStack Swift by following the instructions in Configuring a cloud-container storage pool for data storage.

Encrypting data for cloud-container storage pools

Data that is stored in off-premises cloud-container pools is encrypted by default. You can optionally encrypt data in on-premises cloud-container storage pools.

About this task

For information about encrypting cloud-container storage pool data and for performance considerations related to the encryption of data, see technote 1963635.

Defining a storage rule for cloud tiering

You can define a storage rule to implement cloud tiering, which moves data from a directory-container storage pool on disk to a cloud-container storage pool. The storage rule schedules cloud tiering from directory-container storage pools to cloud-container storage pools.

Before you begin

Restriction: You can configure on-premises or off-premises cloud tiering only on a Microsoft Azure cloud computing system or on a cloud computing system with the Simple Storage Service (S3) protocol.

Review the following information:

- If a cloud-container storage pool is used only for tiering operations (and not for backup operations), the storage pool does not require a local storage directory (cache).
- All extents that are required to rebuild the cloud object are copied to the tier, if the extents are not already there.
- If data is compressed, encrypted, or both in a directory-container storage pool, the data is moved to the cloud-container storage pool in the same format.

About this task

You can define storage rules to specify the following requirements:

- The length of time that data remains in container storage pools on disk before it is moved to cloud storage.
- Whether a storage rule is active or inactive. Storage rules are run daily at a time that is defined in the storage rule.

Procedure

1. On the Operations Center menu bar, click Storage > Tiering Rules.
2. On the Storage Tiering Rules page, click Create Rule.
3. On the Create Rule page, complete the fields and click Create.

Results

When the storage rule is active, the server determines whether the source directory-container pools contain data that is old enough to move. The server moves eligible data to the target cloud-container storage pools.

Reclaiming space in cloud-container storage pools

You can move data from a larger, fragmented cloud container into a smaller, more fully utilized cloud container. In this way, you can help reduce the cost of using object storage for cloud-container storage pools.

Before you begin

Restrictions:

- You can configure a cloud reclamation rule only on a Microsoft Azure cloud computing system or on a cloud computing system with the Simple Storage Service (S3) protocol.
- Your cloud storage provider might charge a fee for the data movement that results from reclamation operations. Before you schedule reclamation operations, use the Operations Center to calculate the effect that different reclamation thresholds might have.

About this task

Fragmentation occurs in cloud-container storage pools when data is deleted or expired. To reclaim space in a cloud-container storage pool, schedule either a daily or an ad hoc reclamation operation.

Procedure

1. In the Operations Center, create a rule to reclaim space by clicking Storage > Rules.

Alternatively, create the rule by using the `DEFINE STGRULE` command with the `ACTIONTYPE=RECLAIM` setting.

2. Optional: Schedule an ad hoc reclamation operation by issuing the `MOVE CONTAINER` command with the default setting `DEFRAG=YES`.

Results

When the storage rule is active, the server determines whether a cloud container reaches its threshold for unused space. If the container space exceeds the threshold that you set, expired data extents are moved to a new, smaller container.

Related reference:

DEFINE STGRULE (Define a storage rule)

DELETE STGRULE (Delete storage rules for storage pools)

MOVE CONTAINER (Move a container)

QUERY STGRULE (Display storage rule information)

Optimizing performance for cloud object storage

You can configure IBM Spectrum Protect™ to temporarily store data in one or more local storage pool directories during data ingestion. The data is then moved from local storage to the cloud. In this way, you can improve data backup and archive performance.

Before you begin

To optimize backup and archive performance, ensure that IBM Spectrum Protect Version 8.1 is installed.

About this task

After you define a storage pool directory, the IBM Spectrum Protect server uses that directory as a temporary landing spot for the data that you are transferring to cloud object storage. The server uses an automated background process to transfer data from local storage in the directory to cloud object storage. You do not need to take any additional steps to start or manage this transfer process. After the server successfully moves the data from local storage to cloud object storage, the server deletes the data from the directory and releases space for more incoming data.

If storage pool directories contain no more free space, backup operations stop prematurely. To avoid this situation, you can allocate more storage pool directories. You can also wait for the data to be automatically removed from the local directories after the data moves to the cloud. The required number of storage pool directories you need to define depends on your disk configuration on the server. When the initial backups occur, the server spreads the data across all the directories you defined.

The amount of space you need for local storage is based on the amount of data you expect to back up each day after data deduplication and compression. If you have a stable network connection to the cloud object storage, the required amount of space is similar to the amount that is required for a daily backup.

For additional planning information, see the topic for your operating system:

- AIX®: Planning for directory-container and cloud-container storage pools
- Linux: Planning for directory-container and cloud-container storage pools
- Windows: Planning for directory-container and cloud-container storage pools

Procedure

1. Create a cloud-container storage pool by using the Add Storage Pool wizard in the Operations Center. Alternatively, create the pool by using the DEFINE STGPOOL command.
2. Define one or more storage pool directories by using the DEFINE STGPOOLDIRECTORY command. Ensure that each storage pool directory has its own file system. On Linux systems, use xfs or ext4 as the file system instead of ext3 because deleting large files takes more time with ext3. Ensure that the new storage pool directories do not share the root file system, nor should they share the same file systems that are used by other IBM Spectrum Protect resources, such as the database or the logs.

Related reference:

DEFINE STGPOOLDIRECTORY (Define a storage pool directory)

AIX

Linux

Windows

Managing space in container storage pools

After you configure IBM Spectrum Protect™ and add storage, manage your data and storage pool space effectively to ensure that it operates correctly. Use container storage pools to maximize your storage space and server performance.

About this task

Container storage pools are primary storage pools that you use for inline data deduplication, inline compression, and cloud storage.

Restriction: You cannot use any of the following functions with container storage pools:

- Migration
- Reclamation
- Aggregation
- Collocation
- Export
- Import
- Simultaneous-write
- Storage pool backup
- Virtual volumes

Procedure

1. Create a directory-container storage pool by completing the following steps:
 - a. Open the Operations Center.
 - b. On the Operations Center menu bar, click Storage > Storage Pools.
 - c. Click +Storage Pool.
 - d. Complete the steps in the Add Storage Pool wizard:
 - To use inline data deduplication, select a Directory storage pool under Container-based storage.
 - When you configure directories for the directory-container storage pool, specify the directory paths that you created for storage during system setup.
 - e. After you configure the new directory-container storage pool, click Close & View Policies to update a management class and start using the storage pool.
2. For optimal performance of container storage pools, complete the following tasks:

Task	Procedure	More information
Protect the storage pool	<p>When you create a directory-container storage pool in the Operations Center, you can configure storage pool protection in the schedule that you assign to the storage pool.</p> <p>Alternatively, use the PROTECT STGPOOL command from the source server to back up data extents in a directory-container storage pool.</p> <p>By protecting a storage pool, you do not use resources that replicate existing data and metadata, which improves server performance.</p>	<ul style="list-style-type: none"> ◦ Protecting data in directory-container storage pools ◦ PROTECT STGPOOL (Protect data that belongs to a storage pool)
Repair a storage pool	<p>When a storage pool is protected, you can use the REPAIR STGPOOL command to repair damaged data extents. Use the REPAIR STGPOOL command to repair a directory-container storage pool.</p> <p>Restriction: If you replicate client nodes but do not protect the directory-container storage pool, you cannot repair the storage pool.</p>	<ul style="list-style-type: none"> ◦ Repairing storage pools ◦ REPAIR STGPOOL (Repair a directory-container storage pool)
Delete containers	<p>Containers are deleted in the inventory as file data is removed or expired.</p> <p>Use the DEFINE STGPOOL command and specify the REUSEDelay parameter to control the duration that</p>	<ul style="list-style-type: none"> ◦ DEFINE STGPOOL (Define a directory-container storage pool) ◦ AUDIT CONTAINER (Verify the consistency of database)

	deduplicated extents are associated with a directory-container storage pool after they are no longer referenced. If a container is damaged, use the AUDIT CONTAINER command to recover or remove data.	information for a directory-container)
Convert a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL)	You can convert an existing storage pool to a directory-container storage pool by completing the steps in Converting a primary storage pool to a container storage pool. Restriction: You cannot convert the following types of storage pool: <ul style="list-style-type: none"> o Primary storage pools that use random-access device classes (DISK) o Copy storage pools o Active-data storage pools 	o CONVERT STGPOOL (Convert a storage pool to a container storage pool)
Monitor container storage pool occupancy	Monitor the storage solution to identify existing and potential issues. For more information, see Monitoring storage solutions.	

- Converting a primary storage pool to a container storage pool
Convert a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL) to a container storage pool. Data that is stored in a container storage pool can use both inline data deduplication and inline compression.
- Cleaning up data in a source storage pool
To convert a storage pool to a directory-container storage pool, you might have to clean up damaged data or files that are in the source storage pool.

AIX | Linux | Windows

Converting a primary storage pool to a container storage pool

Convert a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL) to a container storage pool. Data that is stored in a container storage pool can use both inline data deduplication and inline compression.

Before you begin

To ensure that volumes in a source storage pool and associated copy storage pools are not reused during a conversion process, specify a value for the REUSEDELAY parameter on the UPDATE STGPOOL command. Specify a value for the REUSEDELAY parameter that is greater than the conversion duration. You might have to delay reuse of volumes for the following reasons:

- You inadvertently delete the data during storage pool conversion.
- You need source storage pool functionality that is not available in container storage pools.

Tip: When you specify the REUSEDELAY parameter and a conversion operation is in progress, some storage space is unavailable in the source storage pool until the value of the parameter expires.

Create a container storage pool where data will be moved by completing the following steps:

1. On the Storage Pools page of the Operations Center, click + Storage Pool.
2. Complete the steps in the Add Storage Pool wizard. Select the type of container-based storage that you require.

About this task

By converting a storage pool to a container storage pool, you remove the need for volume reclamation. The omission of volume reclamation operations can help to improve server performance and reduce the amount of required storage hardware.

As files are converted, any copies that are stored in copy pools or active-data pools are deleted.

Restrictions:

- If the source pool is specified as a backup, archive, or migration destination in an active policy set that has pending changes, you must activate those changes before you can convert the pool.
- To ensure that the destination specifies a storage pool that is not converted or undergoing conversion, you must update all policies that reference the source storage pool.
- If the source storage pool is specified as a next storage pool, you must update the NEXTSTGPOOL parameter on the UPDATE STGPOOL command to specify a random-access or sequential-access storage pool that is not being converted.
- The following data types are not eligible for conversion: table of contents (TOC) backups, virtual volumes, and Network Data Management Protocol (NDMP) data. Before you start the conversion process, manually delete these data types from the storage pool, move the data types to a different primary storage pool, or allow the data types to expire based on policy settings.
- When you convert a storage pool with a FILE device class to a directory-container pool, the target storage pool should be approximately 30% larger than the source storage pool. Additional space is not typically required when you convert other storage pool types.

For more information about best practices for storage pool conversion, see [Best practices for IBM Spectrum Protect storage pool conversion](#).

- If the source storage pool is used to store TOC backups, ensure that another primary storage pool is available to store new TOC backups. Existing TOC backups are not moved during conversion.

The TOC pool must use a NATIVE or NONBLOCK data format and a device class other than Centera. To avoid mount delays, use a DISK or FILE device class.

Procedure

1. On the Storage Pools page of the Operations Center, select a storage pool that uses a FILE device class, a tape device class, or VTL.
2. Click More > Convert and complete the steps in the Convert Storage Pool wizard.
Tip: Schedule conversion for at least 2 hours for a storage pool that uses a FILE device class and at least 4 hours for VTL.

What to do next

When the conversion process is complete, the source storage pool might contain damaged data or data that is incompatible with container storage pools. Clean up the source storage pool by completing the steps in [Cleaning up objects after storage pool conversion](#).

Related tasks:

[Restoring the database](#)

AIX

Linux

Windows

Cleaning up data in a source storage pool

To convert a storage pool to a directory-container storage pool, you might have to clean up damaged data or files that are in the source storage pool.

Procedure

Use the following options to recover or repair damaged data:

- Recover an undamaged version of the data from a copy or active-data storage pool by issuing the RESTORE STGPOOL command.
- Recover an undamaged version of the data from a target replication server by issuing the REPLICATE NODE command and specifying the RECOVERDAMAGED=YES parameter.
- Remove data that cannot be repaired after storage pool conversion by issuing the REMOVE DAMAGED command. The REMOVE DAMAGED command might not remove volumes that are marked as destroyed on the source storage pool. To remove these volumes, complete the following steps:
 - a. Issue the DELETE VOLUME command and specify the DISCARDATA=YES parameter.
 - b. Issue the CONVERT STGPOOL command to convert the storage pool again.
 - c. If damaged data is identified during storage pool conversion, reissue the REMOVE DAMAGED command.
- Complete the analysis tasks that are described in [technote 1666371](#).

What to do next

After you have recovered or repaired the damaged data, retry conversion by issuing the CONVERT STGPOOL command. To view information about damaged files that remain in the source storage pool, issue the QUERY CLEANUP command. Tip: If a Cleanup status is shown for a storage pool that contains no data, you can delete the storage pool by using the DELETE STGPOOL command.

Related reference:

DELETE VOLUME (Delete a storage pool volume)
QUERY CLEANUP (Query the cleanup that is required in a source storage pool)
REMOVE DAMAGED (Remove damaged data from a source storage pool)
REPLICATE NODE (Replicate data in file spaces that belong to a client node)
RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)

Auditing a storage pool

You can schedule audit operations to identify corrupted files in storage pools.

Procedure

Issue the DEFINE STGRULE command with the ACTIONTYPE=AUDIT setting.

For more information about the DEFINE STGRULE command, see DEFINE STGRULE (Define a rule for auditing storage pools).

Results

When the storage rule is active, audit operations run according to the defined schedule. You can view information about corrupted files by issuing the QUERY DAMAGED command.

If you detect corrupted files, you can restore data based on your configuration. If you protected the contents of the storage pool by using the PROTECT STGPOOL command, you can repair the contents of the storage pool by using the REPAIR STGPOOL command.

Related reference:

PROTECT STGPOOL (Protect data that belongs to a storage pool)
QUERY DAMAGED (Query damaged in a directory-container or cloud-container storage pool)
REPAIR STGPOOL (Repair a directory-container storage pool)

Auditing a storage pool container

Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.

About this task

You audit a storage pool container in the following situations:

- When you issue the QUERY DAMAGED command and a problem is detected
- When the server displays messages about damaged data extents
- Your hardware reports an issue and error messages that are associated with the storage pool container are displayed

Procedure

1. To audit a storage pool container, issue the AUDIT CONTAINER command. For example, issue the following command to audit a container, 00000000000076c.dcf:

```
audit container c:\tsm-storage\07\00000000000076c.dcf
```

2. Review the output from the ANR4891I message for information about any damaged data extents.

What to do next

If you detect problems with the storage pool container, you can restore data based on your configuration. You can repair the contents in the storage pool by using the REPAIR STGPOOL command.

Restriction: You can repair the contents of the storage pool only if you protected the storage pool by using the PROTECT STGPOOL command.

Related reference:

🔗 [AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

🔗 [QUERY DAMAGED](#) (Query damaged data in a directory-container or cloud-container storage pool)

Storage system requirements and reducing the risk of data corruption

You can use many types of storage for the IBM Spectrum Protect™ server. If you use block disk storage, solid-state drives (SSD), or network-attached file systems for server storage, ensure that the storage meets requirements.

The following requirements apply to storage for the server database, active log, and archive log; for storage pools that use DISK or FILE device classes; and for directory-container storage pools.

Storage can be connected to the server system by any method that is valid for the operating system. For example, the storage can be attached directly, or by using Fibre Channel or iSCSI technology.

Because of the many storage systems that can meet the requirements for server storage, a list of such devices is not available. Contact the vendor if you have questions about whether a system meets IBM Spectrum Protect requirements.

For details about file system requirements, see technote 1902417. For details about network file system (NFS) requirements, see technote 1470193.

Storage and file systems must report write and commit results synchronously and accurately to the IBM Spectrum Protect server. Unreported or asynchronously reported write errors that result in data not being permanently committed to the storage system can cause data corruption. Data corruption can cause operational failures, including failure to start the server, and data recovery is typically required.

You can reduce the risk of data corruption with the following tips:

Write cache

Disk systems use write cache to improve system performance. To reduce the risk of data corruption, the storage system must reliably commit the data in the write cache to permanent storage.

The write cache typically has a battery to prevent loss of data from the cache during short power outages. For critical systems, consider backup power sources to protect the cache from extended power outages.

Direct I/O

Direct I/O meets the server's need for synchronous and accurate reporting on data write and commit operations.

Attention: Do not disable direct I/O in situations where the method of write caching has a potential to cause data loss.

Disabling direct I/O can greatly increase the potential for data loss because more data is cached by the file system, in addition to the disk system.

Storage replication

Environments that replicate IBM Spectrum Protect storage must use features such as maintenance of write order between the source (local server) and the target (remote server). The database, active log, archive logs, and storage pools must be part of a consistency group. A consistency group maintains relationships among volumes to preserve write order so that they can be recovered. Any I/O to the members of the target consistency group must be written in the same order as the source and maintain the same volatility characteristics.

To maintain synchronization between IBM Spectrum Protect servers at local and remote sites, do not start a server at the remote site except in a failover situation. Monitor for synchronization of data at the local and remote locations. If synchronization is lost, you must restore the server at the remote location by using IBM Spectrum Protect restore commands for the database and storage pools.

Tips on storage configuration

For tips on storage configuration to optimize system performance, see the following topics from the V7.1.1 product documentation. The information in the checklists can be applied to later releases.

- Checklist for server database disks
- Checklist for server recovery log disks
- Checklist for storage pools that use DISK or FILE device classes

Monitoring storage solutions

After you implement an IBM Spectrum Protect™ solution, monitor the solution to ensure that it operates correctly. By monitoring the solution on a daily and periodic basis, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

About this task

The preferred way to monitor a solution is by using the Operations Center, which provides overall and detailed system status in a graphical user interface. In addition, you can configure the Operations Center to generate email reports that summarize system status.

Procedure

1. Complete daily monitoring tasks. For instructions, see [Daily monitoring checklist](#).
2. Complete periodic monitoring tasks. For instructions, see [Periodic monitoring checklist](#).
3. To verify that your system complies with licensing requirements, follow the instructions in [Verifying license compliance](#).
4. Optional: Set up email reports of system status. For instructions, see [Tracking system status by using email reports](#)
5. Optional: In some cases, you might want to use advanced monitoring tools to complete specific monitoring or troubleshooting tasks. To select and configure advanced monitoring tools, see [Selecting, configuring, and using monitoring tools](#).

What to do next

To help you diagnose issues with backup-archive clients, install IBM Spectrum Protect client management services on backup-archive client systems that support it. When the client management service is installed on a system, in the Operations Center you can click [Diagnose](#) to get help with diagnosing issues with the backup-archive client. To install the client management service, follow the instructions in [Collecting diagnostic information with IBM Spectrum Protect client management services](#).

Related concepts:

[Performance](#)

Related tasks:


[Troubleshooting](#)


Daily monitoring checklist

Review the checklist to ensure that you complete important daily monitoring tasks.

Complete the daily monitoring tasks from the Operations Center Overview page. You can access the Overview page by opening the Operations Center and clicking [Overviews](#).

The following figure shows the location for completing each task.








Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.


The following table lists the daily monitoring tasks and provides instructions for completing each task.

Table 1. Daily monitoring tasks



Task	Basic procedures	Advanced procedures and troubleshooting information
------	------------------	---



Task	Basic procedures	Advanced procedures and troubleshooting information
<p>Watch for security notifications, which can indicate a ransomware attack.</p>	<p>If a potential ransomware attack is detected in the IBM Spectrum Protect™ environment, a security notification message is displayed in the foreground of the Operations Center. For more information, click the message to open the Security Notifications page.</p>	<p>On the Security Notifications page, you can take the following actions:</p> <ul style="list-style-type: none"> • View notification details by client. Restriction: In Operations Center Version 8.1.5, notifications are available only for backup-archive clients. • Acknowledge a security notification by selecting it and clicking Acknowledge. When you acknowledge a security notification, a check mark is added to the Acknowledged column of the Security Notifications page for the selected client. The standard by which a notification is acknowledged is determined by your organization. A check mark might mean that you investigated the issue and determined that it is a false positive. Or it might mean that a problem exists and is being resolved. • Assign a security notification to an administrator by selecting the security notification and clicking Assign. To view the assignment, the administrator must sign in to the Operations Center and click Overviews > Security. If you are not certain whether the administrator regularly monitors the Security Notifications page, notify the administrator about the assignment. • If the notification is a false positive, you can select the security notification and click Reset. The security notification is deleted. Historical data that is used for baseline comparisons with the most recent backup operation is deleted. A new baseline is calculated going forward.
<p>1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.</p>	<p>To verify whether clients are at risk, in the Clients area, look for an At risk notification. To view details, click the Clients area. Attention: If the At risk percentage is much greater than usual, it might indicate a ransomware attack. A ransomware attack can cause backup operations to fail, thus placing clients at risk. For example, if the percentage of clients at risk is normally between 5% and 10%, but the percentage increases to 40% or 50%, investigate the cause. If you installed the client management service on a backup-archive client, you can view and analyze the client error and schedule logs by completing the following steps:</p> <ol style="list-style-type: none"> 1. In the Clients table, select the client and click Details. 2. To diagnose an issue, click Diagnosis. 	<p>For clients that do not have the client management service installed, access the client system to review the client error logs.</p>


Task	Basic procedures	Advanced procedures and troubleshooting information
<p>2 Determine whether client-related or server-related errors require attention.</p>	<p>To determine the severity of any reported alerts, in the Alerts area, hover over the columns.</p>	<p>To view additional information about alerts, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Alerts area. 2. In the Alerts table, select an alert. 3. In the Activity Log pane, review the messages. The pane displays related messages that were issued before and after the selected alert occurred.
<p>3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.</p>	<ol style="list-style-type: none"> 1. To verify whether servers are at risk, in the Servers area, look for an Unavailable notification. 2. To view additional information, click the Servers area. 3. Select a server in the Servers table and click Details. 	<p>Tip: If you detect an issue that is related to server properties, update the server properties:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a server and click Details. 2. To update server properties, click Properties.
<p>4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.</p>	<ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Status column of the table, view the status of the server and resolve any issues: <ul style="list-style-type: none"> o Normal  Sufficient space is available for the server database, active log, and archive log. o Critical  Insufficient space is available for the server database, active log, or archive log. You must add space immediately, or the data protection services that are provided by the server will be interrupted. o Warning  The server database, active log, or archive log is running out of space. If this condition persists, you must add space. o Unavailable  Status cannot be obtained. Ensure that the server is running, and that there are no network issues. This status is also shown if the monitoring administrator ID is locked or otherwise unavailable on the server. This ID is named IBM-OC-hub_server_name. o Unmonitored  Unmonitored servers are defined to the hub server, but are not configured for management by the Operations Center. To configure an unmonitored server, select the server, and click Monitor Spoke. 	<p>You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>5 Verify server database backup operations.</p>	<p>To determine when a server was most recently backed up, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column. 	<p>To obtain more detailed information about backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. In the Servers table, select a row and click Details. 2. In the DB Backup area, hover over the check marks to review information about backup operations. <p>If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click Back Up. <p>To determine whether the server database is configured for automatic backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the menu bar, hover over the settings icon  and click Command Builder. 2. Issue the QUERY DB command: <pre>query db f=d</pre> 3. In the output, review the Full Device Class Name field. If a device class is specified, the server is configured for automatic database backups.
<p>6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.</p>	<p>To search for information about processes that failed because of server issues, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click Servers > Maintenance. 2. To obtain the two-week history of a process, view the History column. 3. To obtain more information about a scheduled process, hover over the check box that is associated with the process. 	<p>For more information about monitoring processes and resolving issues, see the Operations Center online help.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>7 Verify that the amount of data that was recently sent to and from servers is within the expected range.</p>	<ul style="list-style-type: none"> To obtain an overview of activity in the last 24 hours, view the Activity area. To compare activity in the last 24 hours with activity in the previous 24 hours, review the figures in the Current and Previous areas. 	<ul style="list-style-type: none"> If more data was sent to the server than you expected, determine which clients are backing up more data and investigate the cause. It is possible that client-side data deduplication is not working correctly. Attention: If the amount of backed-up data is significantly larger than usual, it might indicate a ransomware attack. When ransomware encrypts data, the system perceives the data as being changed, and the changed data is backed up. Thus, backup volumes become larger. To determine which clients are affected, click the Applications, Virtual, or Systems tabs. If less data was sent to the server than you expected, investigate whether client backup operations are proceeding on schedule.

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>8 Verify that storage pools are available to back up client data.</p>	<ol style="list-style-type: none"> 1. If problems are indicated in the Storage & Data Availability area, click Pools to view the details: <ul style="list-style-type: none"> o If the Critical  status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. Attention: If the status is critical, investigate the cause: <ul style="list-style-type: none"> ■ If the data deduplication rate for a storage pool drops significantly, it might indicate a ransomware attack. During a ransomware attack, data is encrypted and cannot be deduplicated. To verify the data deduplication rate, in the Storage Pools table, review the value in the % Savings column. ■ If a storage pool unexpectedly becomes 100% utilized, it might indicate a ransomware attack. To verify the utilization, review the value in the Capacity Used column. Hover over the values to see the percentages of used and free space. o If the Warning  status is displayed, the storage pool is running out of space, or its access status is read-only. 2. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column. 	<p>To view the storage-pool capacity that was used over the past two weeks, select a row in the Storage Pools table and click Details.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>9 Verify that storage devices are available for backup operations.</p>	<p>In the Storage & Data Availability area, in the Volumes section, under the capacity bars, review the status that is reported next to Devices. If a Critical  or Warning  status is displayed for any device, investigate the issue. To view details, click Devices.</p>	<p>Disk devices might have a critical or warning status for the following reasons:</p> <ul style="list-style-type: none"> • For DISK device classes, volumes might be offline or have a read-only access status. The Disk Storage column of the Disk Devices table shows the state of volumes. • For FILE device classes that are not shared, directories might be offline. Also, insufficient free space might be available for allocating scratch volumes. The Disk Storage column of the Disk Devices table shows the state of directories. • For FILE device classes that are shared, drives might be unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. Other columns of the Disk Devices table show the state of the drives and paths. <p>Tape devices might have a warning or critical status if drives are unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. A tape device might also have a critical status if the library is offline. Other columns of the Tape Devices table show the state of the library robotics, drives, and paths.</p> <p>For tape backup operations, verify that sufficient scratch tapes are available. If you are not certain whether the number of available scratch tapes is sufficient, open the details notebook to view tape usage and an estimate of scratch tape availability. To open the details notebook, select a library in the table and click Details.</p>

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>10 Monitor node replication processes.</p>	<ol style="list-style-type: none"> To obtain the overall status of node replication processes, view the Replication area on the Operations Center Overview page. To view information about each replicated server pair, click the Replication area. Attention: If you notice an unexpected increase in the number of replication failures, it might indicate a ransomware attack. Investigate the cause of the failures. To view the amount of data that was replicated over the last two weeks and the speed of replication, select a server pair and click Details. To view replication information for a client, on the Operations Center Overview page, click Clients. View the information in the Replication Workload column. Attention: If you see a drastic, unexpected increase in the replication workload, it might indicate a ransomware attack. Investigate the cause of the increased workload. 	<p>For advanced monitoring, view information about running and ended node replication processes by using commands:</p> <ol style="list-style-type: none"> On the Operations Center Overview page, hover over the settings icon  and click Command Builder. Issue the QUERY REPLICATION command. For instructions, see QUERY REPLICATION (Query node replication processes). If the replication operation was completed successfully, the <code>Total Files To Replicate</code> value matches the <code>Total Files Replicated</code> value. <p>To display messages that are related to a node replication process on a source or target replication server, complete the following steps:</p> <ol style="list-style-type: none"> On the Operations Center Overview page, click Servers. Select the source or target replication server and click Details: <ul style="list-style-type: none"> To view active tasks, click Active Tasks, select the task, and verify that the Running status is displayed. For details, view the related activity logs. To view completed tasks, click Completed Tasks, select the task, and ensure that the Completed status is displayed. For details, view the related activity logs.

Periodic monitoring checklist

To help ensure operations run correctly, complete the tasks in the periodic monitoring checklist. Schedule periodic tasks frequently enough so that you can detect potential issues before they become problematic.




Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the settings icon  and click Command Builder.

Table 1. Periodic monitoring tasks


Task	Basic procedures	Advanced procedures and troubleshooting
------	------------------	---

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor system performance.	<p>Determine the length of time that is required for client backup operations:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. Find the server that is associated with the client. 2. Click Servers. Select the server and click Details. 3. To view the duration of completed tasks in the last 24 hours, click Completed Tasks. 4. To view the duration of tasks that were completed more than 24 hours ago, use the QUERY ACTLOG command. Follow the instructions in . 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. <p>If you installed the client management service on a backup-archive client, you can diagnose performance issues for the backup-archive client by completing the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Clients. 2. Select a backup-archive client and click Details. 3. To retrieve client logs, click Diagnosis. 	<p>For instructions about reducing the time that it takes for the client to back up data to the server, see Resolving common client performance problems.</p> <p>Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks.</p> <p>For information about identifying and resolving other performance issues, see Performance.</p>
Determine the disk savings that are provided by data deduplication.	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Pools. 2. Select a pool and click Quick Look. 3. In the Data Deduplication area, view the Space saved row. 	<p>For advanced monitoring, to obtain detailed statistics about the data-deduplication process for a specific directory-container storage pool or cloud-container storage pool, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. Obtain a statistical report by issuing the GENERATE DEDUPSTATS command. Follow the instructions in GENERATE DEDUPSTATS (Generate data deduplication statistics for a directory-container storage pool). 3. View the statistical report by issuing the QUERY DEDUPSTATS command. Follow the instructions in QUERY DEDUPSTATS (Query data deduplication statistics).

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Verify that current backup files for device configuration and volume history information are saved.</p>	<p>Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations.</p> <p>To locate the volume history and device configuration files, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To locate the volume history and device configuration files, issue the following commands: <ul style="list-style-type: none"> <code>query option volhistory</code> <code>query option devconfig</code> 3. In the output, review the Option Setting column to find the file locations. <p>If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.</p>	

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Determine whether sufficient space is available for the instance directory file system.</p>	<p>Verify that at least 20% of free space is available in the instance directory file system. Take the action that is appropriate for your operating system:</p> <ul style="list-style-type: none"> <p>AIX To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -g instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Linux To view available space in the file system, on the operating system command line, issue the following command:</p> <pre>df -h instance_directory</pre> <p>where <i>instance_directory</i> specifies the instance directory.</p> <p>Windows In the Windows Explorer program, right-click the file system and click Properties. View the capacity information.</p> <p>The preferred location of the instance directory depends on the operating system where the server is installed:</p> <ul style="list-style-type: none"> <p>AIX Linux</p> <pre>/home/tsminst1/tsminst1</pre> <p>Windows <pre>C:\tsminst1</pre></p> <p>Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.</p>	
<p>Identify unexpected client activity.</p>	<p>To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps:</p> <ol style="list-style-type: none"> On the Operations Center Overview page, click the Clients area. To view activity over the past two weeks, double-click any client. To view the number of bytes sent to the client, click the Properties tab. In the Last Session area, view the Sent to client row. 	<p>When you double-click a client in the Clients table, the Activity over 2 Weeks area displays the amount of data that the client sent to the server each day.</p> <p>Periodically review the SQL activity summary table, which contains statistics about client sessions. To compare current activity with previous activity, use an SQL SELECT statement. If the level of activity is significantly different from previous activity, it might indicate a ransomware attack.</p> <p>Periodically review the activity log. Look for ANE messages that indicate how many files were backed up and inspected. Compare current data deduplication rates with previous rates. If an unusually high number of files were backed up, or the rate of data deduplication unexpectedly drops to 0, it might indicate a ransomware attack.</p>

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Monitor storage pool growth over time.</p>	<ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Pools area. 2. To view the capacity that was used over the last two weeks, select a pool and click Details. 	<p>Tips:</p> <ul style="list-style-type: none"> • To specify the time period that must elapse before all deduplicated extents are removed from a directory-container storage pool or cloud-container storage pool after they are no longer referenced by the inventory, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. 3. Specify the duration in the <code>Delay period for container reuse</code> field. • To determine data deduplication performance for directory-container and cloud-container storage pools, use the <code>GENERATE DEDUPSTATS</code> command. • To view data deduplication statistics for a storage pool, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. <p>Alternatively, use the <code>QUERY EXTENTUPDATES</code> command to display information about updates to data extents in directory-container or cloud-container storage pools. The command output can help you determine which data extents are no longer referenced and which ones are eligible to be deleted from the system. In the output, monitor the number of data extents that are eligible to be deleted from the system. This metric has a direct correlation to the amount of free space that will be available within the container storage pool.</p> <ul style="list-style-type: none"> • To display the amount of physical space that is occupied by a file space after the removal of the data deduplication savings, use the <code>select * from occupancy</code> command. The command output includes the <code>LOGICAL_MB</code> value. <code>LOGICAL_MB</code> is the amount of space that is used by the file space.
<p>Evaluate the timing of client schedules. Ensure that the start and end times of client schedules meet your business needs.</p>	<p>On the Operations Center Overview page, click Clients > Schedules.</p> <p>In the Schedules table, the Start column displays the configured start time for the scheduled operation. To see when the most recent operation was started, hover over the clock icon.</p>	<p>Tip: You can receive a warning message if a client operation runs longer than expected. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over Clients and click Schedules. 2. Select a schedule and click Details. 3. View the details of a schedule by clicking the blue arrow next to the row. 4. In the Run time alert field, specify the time when a warning message will be issued if the scheduled operation is not completed. 5. Click Save.

Task	Basic procedures	Advanced procedures and troubleshooting
Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks meet your business needs.	<p>On the Operations Center Overview page, click Servers > Maintenance.</p> <p>In the Maintenance table, review the information in the Last Run Time column. To see when the last maintenance task was started, hover over the clock icon.</p>	<p>Tip: If a maintenance task is running too long, change the start time or the maximum run time. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To change the start time or maximum run time for a task, issue the UPDATE SCHEDULE command. For instructions, see UPDATE SCHEDULE (Update a client schedule).

Related reference:

- [UPDATE STGPOOL](#) (Update a storage pool)
- [QUERY EXTENTUPDATES](#) (Query updated data extents)

Verifying license compliance

Verify that your IBM Spectrum Protect™ solution complies with the provisions of your licensing agreement. By verifying compliance regularly, you can track trends in data growth or processor value unit (PVU) usage. Use this information to plan for future license purchasing.

About this task

The method that you use to verify that your solution complies with license terms varies depending on the provisions of your IBM Spectrum Protect licensing agreement.

Front-end capacity licensing

The front-end model determines license requirements based on the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems.

Back-end capacity licensing

The back-end model determines license requirements based on the terabytes of data that are stored in primary storage pools and repositories.

Tips:

- To ensure the accuracy of front-end and back-end capacity estimates, install the most recent version of the client software on each client node.
- The front-end and back-end capacity information in the Operations Center is for planning and estimation purposes.

PVU licensing

The PVU model is based on the use of PVUs by server devices.

Important: The PVU calculations that are provided by IBM Spectrum Protect are considered estimates and are not legally binding. The PVU licensing information that is reported by IBM Spectrum Protect is not considered an acceptable substitute for the IBM® License Metric Tool.



For the most recent information about licensing models, see the information about product details and licenses at the IBM Spectrum Protect product family website. If you have questions or concerns about licensing requirements, contact your IBM Spectrum Protect software provider.

Procedure

To monitor license compliance, complete the steps that correspond to the provisions of your licensing agreement.

Tip: The Operations Center provides an email report that summarizes front-end and back-end capacity usage. Reports can be sent automatically to one or more recipients regularly. To configure and manage email reports, click Reports on the Operations Center menu bar.

Option	Description
--------	-------------

Option	Description
Front-end model	<p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>The front-end capacity estimate is displayed on the Front-end Usage page.</p> <p>b. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage.</p> <p>c. To estimate capacity for clients that did not report capacity usage, go to the following FTP site, which provides measuring tools and instructions:</p> <p><code>ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools</code></p> <p>To measure front-end capacity by script, complete the instructions in the most recently available licensing guide.</p> <p>d. Add the Operations Center estimate and any estimates that you obtained by using a script.</p> <p>e. Verify that the estimated capacity complies with your licensing agreement.</p>
Back-end model	<p>Restriction: If the source and target replication servers do not use the same policy settings, you cannot use the Operations Center to monitor back-end capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>b. Click the Back-end tab.</p> <p>c. Verify that the estimated amount of data complies with your licensing agreement.</p>
PVU model	<p>For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model.</p>

- Assessing compliance with the PVU licensing model
If you purchased IBM Spectrum Protect under the processor value unit (PVU) licensing model, ensure that your solution complies with the license terms. Review the PVU estimates periodically to plan for future license purchasing. For example, if PVU estimates increase or you plan to install more servers, you might have to purchase more licenses.

Tracking system status by using email reports

Set up the Operations Center to generate email reports that summarize system status. You can configure a mail server connection, change report settings, and optionally create custom reports.

Before you begin

Before you set up email reports, ensure that the following requirements are met:

- A Simple Mail Transfer Protocol (SMTP) host server is available to send and receive reports by email. The SMTP server must be configured as an open mail relay. You must also ensure that the IBM Spectrum Protect™ server that sends email messages has access to the SMTP server. If the Operations Center is installed on a separate computer, that computer does not require access to the SMTP server.
- To set up email reports, you must have system privilege for the server.
- To specify the recipients, you can enter one or more email addresses or administrator IDs. If you plan to enter an administrator ID, the ID must be registered on the hub server and must have an email address that is associated with it. To specify an email address for an administrator, use the EMAILADDRESS parameter of the UPDATE ADMIN command.

About this task

You can configure the Operations Center to send a general operations report, a license compliance report, and one or more custom reports. You create custom reports by selecting a template from a set of commonly used report templates or by entering SQL SELECT statements to query managed servers.

Procedure

To set up and manage email reports, complete the following steps:

1. On the Operations Center menu bar, click Reports.
2. If an email server connection is not yet configured, click Configure Mail Server and complete the fields. After you configure the mail server, the general operations report and license compliance report are enabled.
3. To change report settings, select a report, click Details, and update the form.
4. Optional: To add a custom report, click + Report, and complete the fields.
Tip: To immediately run and send a report, select the report and click Send.

Results

Enabled reports are sent according to the specified settings.

Related reference:

[UPDATE ADMIN](#) (Update an administrator)

Related information:

[Custom Report Examples](#)

Selecting, configuring, and using monitoring tools

Use the Operations Center to obtain an overview of system status and to drill down to more detailed information. In some cases, you might want to use advanced tools to collect specific monitoring information.

Procedure

Select and configure the monitoring tools that are appropriate for your solution.

Table 1. Monitoring tools

Tool type	Use cases	Links to more information
Operations Center	<ul style="list-style-type: none">• Use the graphical user interface to review system status and diagnose issues.• Set up the Operations Center to send daily email summary reports.• Optional: Customize the alerts that are displayed in the Operations Center and set up email notifications about the alerts.• Optional: Monitor the storage environment remotely by viewing the Overview page in the web browser of a mobile device. For example, you can use the Apple Safari web browser on an Apple iPad device. Other mobile devices can also be used. <p>Tip: If you install IBM Spectrum Protect™ client management services on a backup-archive client, you can use the Operations Center to obtain troubleshooting information for the backup-archive client. The client management service can be installed only on Linux or Windows operating systems.</p>	

Tool type	Use cases	Links to more information
IBM Spectrum Protect administrative commands	<p>Review detailed information. Use the method that is appropriate for your solution:</p> <ul style="list-style-type: none"> To display messages that were generated by the server and client, use the QUERY ACTLOG command. Tip: You can run administrative commands from the Operations Center command builder. To monitor activities such as server migration and client logons, use the administrative client in console mode. Run the dsmadm -consolemode command. 	<ul style="list-style-type: none"> Administrative commands QUERY ACTLOG (Query the activity log) Monitoring server activities from the administrative client Administrative client options
Event logging	Log server messages and most client messages as events to one or more repositories called receivers.	<p>AIX Linux Windows For instructions about using event logging to monitor a solution, see Logging IBM Spectrum Protect events to receivers (V7.1.1).</p> <p>Linux For instructions about logging events to a Linux system log, see Logging events to the Linux system log (V7.1.4).</p>
SQL queries	<p>Create and format customized queries of the server database. For example, you can query the SQL activity summary table to view statistics about client operations and server processes. To display all information in the summary table, issue the following command from the administrative client:</p> <pre>select * from summary</pre>	Using SELECT commands (V7.1.1)
Operating system tools	Monitor and test system performance.	
Device-monitoring tools	Monitor devices for availability, capacity, and performance. For example, use IBM Spectrum Control™ or tools that are included in device hardware packages.	<p>To monitor overall device status by using IBM Spectrum Control, follow the instructions in Monitoring the status and condition of resources.</p> <p>To monitor performance by using IBM Spectrum Control, follow the instructions in Monitoring the performance of resources.</p>

Tool type	Use cases	Links to more information
IBM® Tivoli® Monitoring for Tivoli Storage Manager	Monitor IBM Spectrum Protect servers and produce historical reports about server and client activities. Tip: The Operations Center is the preferred tool for monitoring. However, Tivoli Monitoring for Tivoli Storage Manager is useful for generating historical reports that are based on IBM Cognos® Business Intelligence technology.	Tivoli Monitoring for Tivoli Storage Manager

Managing operations

By effectively managing server and client operations, you can optimize the performance of your storage environment. To get started, monitor the environment by using the Operations Center. Then, take action to prevent potential issues and improve performance.

About this task

- **Managing server operations**
You can start and stop the server, manage inventory capacity, and manage memory and processor usage. You can also optimize data transfer between servers, upgrade the server, and tune scheduled activities.
- **Managing client operations**
You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.
- **Managing the Operations Center**
The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect™ environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

Managing server operations

You can start and stop the server, manage inventory capacity, and manage memory and processor usage. You can also optimize data transfer between servers, upgrade the server, and tune scheduled activities.

- **Stopping and starting the server**
Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.
- **Managing inventory capacity**
Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.
- **Managing memory and processor usage**
Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.
- **Determining whether Aspera FASP technology can optimize data transfer in your system environment**
If your IBM Spectrum Protect™ server replicates nodes or protects storage pools to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Before you enable Aspera FASP technology, you must obtain the appropriate licenses. Both evaluation and full licenses are available.
- **Planning to upgrade the server**
When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.
- **Tuning scheduled activities**
Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.

Before you begin

You must have system or operator privilege to stop and start the IBM Spectrum Protect™ server.

- Stopping the server
Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.
- Starting the server for maintenance or reconfiguration tasks
Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Stopping the server

Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.

About this task

When you issue the HALT command to stop the server, the following actions occur:

- All processes and client node sessions are canceled.
- All current transactions are stopped. (The transactions will be rolled back when the server is restarted.)

Procedure

To prepare the system and stop the server, complete the following steps:

1. Prevent new client node sessions from starting by issuing the DISABLE SESSIONS command:

```
disable sessions all
```

2. Determine whether any client node sessions or processes are in progress by completing the following steps:
 - a. On the Overview page of the Operations Center, view the Activity area for the total numbers of processes and sessions that are currently active. If numbers differ significantly from the usual numbers that are displayed during your daily storage-management routine, view other status indicators in the Operations Center to check whether there is a problem.
 - b. View the graph in the Activity area to compare the amount of network traffic over the following periods:
 - The current period, that is, the most recent 24-hour period
 - The previous period, that is, the 24 hours before the current periodIf the graph for the previous period represents the expected amount of traffic, significant differences on the graph for the current period might indicate a problem.
 - c. On the Servers page, select a server for which you want to view processes and sessions, and click Details. If the server is not registered as a hub or spoke server in the Operations Center, obtain information about processes by using administrative commands. Issue the QUERY PROCESS command to query processes and obtain information about sessions by issuing the QUERY SESSION command.
3. Wait until the client node sessions are completed or cancel them. To cancel processes and sessions, complete the following steps:
 - On the Servers page, select a server for which you want to view processes and sessions, and click Details.
 - Click the Active Tasks tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
 - Click Cancel.
 - If the server is not registered as a hub or spoke server in the Operations Center, cancel sessions by using administrative commands. Issue the CANCEL SESSION command to cancel a session and cancel processes by using the CANCEL PROCESS command.

Tip: If the process that you want to cancel is waiting for a tape volume to be mounted, the mount request is canceled. For example, if you issue an EXPORT, IMPORT, or MOVE DATA command, the command might initiate a

process that requires a tape volume to be mounted. However, if a tape volume is being mounted by an automated library, the cancel operation might not take effect until the mount process is complete. Depending on your system environment, this could take several minutes.

4. Stop the server by issuing the HALT command:

```
halt
```

Starting the server for maintenance or reconfiguration tasks

Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

About this task

Start the server in maintenance mode by running the DSMSEV utility with the MAINTENANCE parameter.

The following operations are disabled in maintenance mode:

- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the HALT command:

```
halt
```

2. Start the server by using the method that you use in production mode. Follow the instructions for your operating system:

- o **AIX** Starting the server instance
- o **Linux** Starting the server instance
- o **Windows** Starting the server instance

Operations that were disabled during maintenance mode are reenabled.

Managing inventory capacity

Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.

Before you begin

The active and archive logs have the following characteristics:

- The active log can be a maximum size of 512 GB. For more information about sizing the active log for your system, see Planning the storage arrays.
- The archive log size is limited to the size of the file system that it is installed on. The archive log size is not maintained at a predefined size like the active log. Archive log files are automatically deleted after they are no longer needed.

As a best practice, you can optionally create an archive failover log to store archive log files when the archive log directory is full.

Check the Operations Center to determine the component of the inventory that is full. Ensure that you stop the server before you increase the size of one of the inventory components.

Procedure

- To increase the size of the database, complete the following steps:
 - Create one or more directories for the database on separate drives or file systems.
 - Issue the EXTEND DBSPACE command to add the directory or directories to the database. The directories must be accessible to the instance user ID of the database manager. By default, data is redistributed across all database directories and space is reclaimed.

Tips:

 - The time that is needed to complete redistribution of data and reclaiming of space is variable, depending on the size of your database. Make sure that you plan adequately.
 - Ensure that the directories that you specify are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.
 - Halt and restart the server to fully use the new directories.
 - Reorganize the database if necessary. Index and table reorganization for the server database can help to avoid unexpected database growth and performance issues. For more information about reorganizing the database, see technote 1683633.

- To decrease the size of the database for V7.1 servers and later, issue the following DB2® commands from the server instance directory:

Restriction: The commands can increase I/O activity, and might affect server performance. To minimize performance problems, wait until one command is completed before you issue the next command. The DB2 commands can be issued when the server is running.

```
db2 connect to tsmdb1
db2 set schema tsmdb1
db2 ALTER TABLESPACE USERSPACE1 REDUCE MAX
db2 ALTER TABLESPACE IDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGEIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGESPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE5 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIIDXSPACE5 REDUCE MAX
```

- To increase or decrease the size of the active log, complete the following steps:
 1. Ensure that the location for the active log has enough space for the increased log size. If a log mirror exists, its location must also have enough space for the increased log size.
 2. Halt the server.
 3. In the dsmserv.opt file, update the ACTIVELOGSIZE option to the new size of the active log, in megabytes. The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

To change the active log to its maximum size of 512 GB, enter the following server option:

```
activelogsiz 524288
```

4. If you plan to use a new active log directory, update the directory name that is specified in the ACTIVELOGDIRECTORY server option. The new directory must be empty and must be accessible to the user ID of the database manager.
 5. Restart the server.
- Compress the archive logs to reduce the amount of space that is required for storage. Enable dynamic compression of the archive log by issuing the following command:

```
setopt archlogcompress yes
```

Restriction: Use caution when you enable the ARCHLOGCOMPRESS server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log directory file system usage nears out of space conditions, the ARCHLOGCOMPRESS server option must be disabled. You can use the SETOPT command to disable archive log compression immediately without halting the server.

Related reference:

- [ACTIVELOGSIZE server option](#)
- [EXTEND DBSPACE \(Increase space for the database\)](#)
- [SETOPT \(Set a server option for dynamic update\)](#)

Managing memory and processor usage

Ensure that you manage memory requirements and processor usage to ensure that the server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.

Before you begin

- Ensure that your configuration uses the required hardware and software. For more information, see IBM Spectrum Protect™ Supported Operating Systems.
- For more information about managing resources such as the database and recovery log, see Planning the storage arrays.
- Add more system memory to determine whether there is a performance improvement. Monitor memory usage regularly to determine whether more memory is required.

Procedure

1. Release memory from the file system cache where possible.
2. To manage the system memory that is used by each server on a system, use the DBMEMPERCENT server option. Limit the percentage of system memory that can be used by the database manager of each server. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.
3. Set the user data limit and private memory for the database to ensure that private memory is not exhausted. Exhausting private memory can result in errors, less than optimal performance, and instability.

Linux

Determining whether Aspera FASP technology can optimize data transfer in your system environment

If your IBM Spectrum Protect™ server replicates nodes or protects storage pools to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Before you enable Aspera FASP technology, you must obtain the appropriate licenses. Both evaluation and full licenses are available.

Before you begin

Aspera FASP technology is used to transfer data extents from a container storage pool to a remote server. When Aspera FASP technology is enabled, the data extents are always encrypted during transfer regardless of whether the Secure Sockets Layer (SSL) protocol is enabled. However, if you want to secure the network connection, enable SSL. For information about SSL and how to enable it, see Secure Sockets Layer and Transport Layer Security communication.

About this task

Restrictions:

- Use Aspera FASP technology when your wide area network (WAN) shows signs of high packet loss, data transfer delays that are caused by network impairment, or both. If WAN performance meets your business needs, do not enable Aspera FASP technology.
- To enable Aspera FASP technology for node replication operations, the data must be stored in a directory-container storage pool.

Procedure

1. Determine whether Aspera FASP technology is appropriate for your system environment. If either of the following conditions occurs, enable Aspera FASP technology:
 - Average delays for data-transfer operations exceed 50 milliseconds.
 - Packet loss is greater than 0.01%.

Network characteristics can vary widely. You might be able to improve network throughput by enabling Aspera FASP technology even if the data-transfer delay is less than 50 milliseconds and the packet loss is less than 0.01%.

2. Obtain and install the appropriate licenses. Take one of the following actions:

Obtain and install evaluation licenses

To obtain and install evaluation licenses, which expire in 30 days, complete the following steps:

- a. Request the licenses by sending an email to alliances@asperasoft.com:
 - Include your company name, address, phone number, and the email address of the primary contact at your company.
 - State that you require a 30-day evaluation license.
 - Indicate the number of licenses that you require.

One license is required for each server that is used for data transfer with Aspera FASP technology. For example, if you are replicating a node from a source server to a target server, you require two licenses.

If the license request is approved, the primary contact can expect to receive an email within 24 hours. The email will have license file attachments that are named according to the following convention:

```
xxxxx-ConnectSrv-unlim.eval.aspera-license
```

where xxxxx is a unique number.

- b. Copy one of the license files to the server bin directory of the source server. Select either license file. By default, the directory is in the following location:

```
/opt/tivoli/tsm/server/bin
```

- c. Copy the remaining license file to the bin directory of the target server.
- d. On the source and target servers, set the permission level of each license file to 755. For example, if you are using the default installation directory and the unique license number is 47474, issue the following command on one line:

```
chmod 755 /opt/tivoli/tsm/server/bin/  
47474-ConnectSrv-unlim.eval.aspera-license
```

Obtain and install full licenses

To obtain and install full, unlimited licenses, which do not expire, complete the following steps:

- a. Purchase the IBM Spectrum Protect High Speed Data Transfer product. The product identification number is 5725-Z10. You can obtain the product from Passport Advantage®.

One instance of IBM Spectrum Protect High Speed Data Transfer is required for each server that is used to transfer data with Aspera FASP technology. For example, if you are replicating a node from a source server to a target server, you require two instances of IBM Spectrum Protect High Speed Data Transfer.

- b. Install IBM Spectrum Protect High Speed Data Transfer on each server by using the installation wizard.

Restriction: If the required licenses are missing or expired, operations to replicate nodes and protect storage pools by using Aspera FASP technology fail.

3. Optional: Validate the Aspera FASP configuration by issuing the VALIDATE ASPERA command. You can use the VALIDATE ASPERA command to verify that your system environment is correctly configured for Aspera FASP and to verify that valid licenses are installed. In addition, you can use the command to compare the speed of network throughput with Aspera FASP and TCP/IP technology.

What to do next

To enable Aspera FASP technology, follow the steps in [Optimizing data transfer by enabling Aspera FASP technology](#).

- [Optimizing data transfer by enabling Aspera FASP technology](#)
If you use a remote server for storage pool protection or node replication and you experience network issues, you might want to optimize data transfer by using Aspera Fast Adaptive Secure Protocol (FASP) technology.

Planning to upgrade the server

When a fix pack or interim fix becomes available, you can upgrade the IBM Spectrum Protect™ server to take advantage of product improvements. Servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.

About this task

Follow these guidelines:

- The preferred method is to upgrade the server by using the installation wizard. After you start the wizard, in the IBM Installation Manager window, click the Update icon; do not click the Install or Modify icon.
- If upgrades are available for both the server component and the Operations Center component, select the check boxes to upgrade both components.

Procedure

1. Review the list of fix packs and interim fixes. See [technote 1239415](#).
2. Review product improvements, which are described in readme files.
Tip: When you obtain the installation package file from the IBM Spectrum Protect support site, you can also access the readme file.
3. Ensure that the version that you upgrade your server to is compatible with other components, such as storage agents and library clients. See [technote 1302789](#).
4. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted. See [technote 1053218](#).
5. Review the upgrade instructions. Ensure that you back up the server database, the device configuration information, and the volume history file.

What to do next

To install a fix pack or interim fix, follow the instructions for your operating system:

- **AIX** [Installing an IBM Spectrum Protect server fix pack](#)
- **Linux** [Installing an IBM Spectrum Protect server fix pack](#)
- **Windows** [Installing an IBM Spectrum Protect server fix pack](#)

Tuning scheduled activities

Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize server resources and effectively use different functions available within your solution.

Procedure

1. Monitor system performance regularly to ensure that client backup and server maintenance tasks are completing successfully. Follow the instructions in [Monitoring storage solutions](#).
2. Optional: If the monitoring information shows that the server workload increased, review the planning information. Review whether the capacity of the system is adequate in the following cases:
 - o The number of clients increases
 - o The amount of data that is being backed up increases
 - o The amount of time that is available for backups changes
3. Determine whether your solution is performing at the level you expect. Review the client schedules to check whether tasks are completing within the scheduled time frame:
 - a. On the Clients page of the Operations Center, select the client.
 - b. Click Details.
 - c. From the client Summary page, review the Backed up and Replicated activity to identify any risks.

Adjust the time and frequency of client backup operations, if necessary.

4. Schedule enough time for the following maintenance tasks to complete successfully within a 24-hour period:
 - a. Protect storage pools.
 - b. Replicate node data.
 - c. Back up the database.
 - d. Run expiration processing to remove client backups and archive file copies from server storage.

Tip: Schedule maintenance tasks to start at an appropriate time and in the correct sequence. For example, schedule replication tasks after client backups complete successfully.

- Moving clients from one server to another
To avoid running out of space on a server or to resolve workload issues, you might have to move client nodes from one server to another.

Related concepts:

[Performance](#)

Related tasks:

[Deduplicating data \(V7.1.1\)](#)

Managing client operations

You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.

About this task

In some cases, you can resolve client errors by stopping and starting the client acceptor. If client nodes or administrator IDs are locked, you can resolve the issue by unlocking the client node or administrator ID, and then resetting the password.

For detailed instructions about identifying and resolving client errors, see [Resolving client problems](#).

- Modifying the scope of a client backup
When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.
- Evaluating errors in client error logs
You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.
- Stopping and restarting the client acceptor
If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.
- Resetting passwords
If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to

access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

- Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect™ server, but the workstation is no longer used, you can decommission the workstation.

- Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Spectrum Protect server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

- Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Modifying the scope of a client backup

When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

About this task

When you exclude unnecessary objects from backup operations, you get better control of the amount of storage space that is required for backup operations, and the cost of storage. Depending on your licensing package, you also might be able to limit licensing costs.

Procedure

How you modify the scope of backup operations depends on the product that is installed on the client node:

- For a backup-archive client, you can create an include-exclude list to include or exclude a file, groups of files, or directories from backup operations. To create an include-exclude list, follow the instructions in [Creating an include-exclude list](#).

To ensure consistent use of an include-exclude list for all clients of one type, you can create a client option set on the server that contains the required options. Then, you assign the client option set to each of the clients of the same type. For details, see [Controlling client operations through client option sets](#).

- For a backup-archive client, you can specify the objects to include in an incremental backup operation by using the domain option. Follow the instructions in [Domain option](#).
- For other products, to define which objects are included in and excluded from backup operations, follow the instructions in the product documentation.

Evaluating errors in client error logs

You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

Before you begin

To resolve errors in a backup-archive client on a Linux or Windows operating system, ensure that the client management service is installed and started. For installation instructions, see [Collecting diagnostic information with client management services](#).

Procedure

To diagnose and resolve client errors, take one of the following actions:

- If the client management service is installed on the client node, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click Details.
 3. On the client Summary page, click the Diagnosis tab.
 4. Review the retrieved log messages.

Tips:

- To show or hide the Client Logs pane, double-click the Client Logs bar.

- To resize the Client Logs pane, click and drag the Client Logs bar.

If suggestions are displayed on the Diagnosis page, select a suggestion. In the Client Logs pane, client log messages to which the suggestion relates are highlighted.

5. Use the suggestions to resolve the problems that are indicated by the error messages.

Tip: Suggestions are provided for only a subset of client messages.

- If the client management service is not installed on the client node, review the error logs for the installed client.

Stopping and restarting the client acceptor

If you change the configuration of your solution, you must restart the client acceptor on all client nodes where a backup-archive client is installed.

About this task

In some cases, you can resolve client scheduling problems by stopping and restarting the client acceptor. The client acceptor must be running to ensure that scheduled operations can occur on the client. For example, if you change the IP address or domain name of the server, you must restart the client acceptor.

Procedure

Follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

- To stop the client acceptor, complete the following steps:
 - a. Determine the process ID for the client acceptor by issuing the following command on the command line:

```
ps -ef | grep dsmcad
```

Review the output. In the following sample output, 6764 is the process ID for the client acceptor:

```
root 6764 1 0 16:26:35 ? 0:00 /usr/bin/dsmcad
```

- b. Issue the following command on the command line:

```
kill -9 PID
```

where *PID* specifies the process ID for the client acceptor.

- To start the client acceptor, issue the following command on the command line:

```
/usr/bin/dsmcad
```

Linux

- To stop the client acceptor (and not restart it), issue the following command:

```
# service dsmcad stop
```

- To stop and restart the client acceptor, issue the following command:

```
# service dsmcad restart
```

MAC OS X

Click Applications > Utilities > Terminal.

- To stop the client acceptor, issue the following command:

```
/bin/launchctl unload -w com.ibm.tivoli.dsmcad
```

- To start the client acceptor, issue the following command:

```
/bin/launchctl load -w com.ibm.tivoli.dsmcad
```

Windows

- To stop the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Stop and OK.
- To restart the client acceptor service, complete the following steps:
 - a. Click Start > Administrative Tools > Services.
 - b. Double-click the client acceptor service.
 - c. Click Start and OK.

Related reference:

[Resolving client scheduling problems](#)

Resetting passwords

If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

Procedure

To resolve password issues, take one of the following actions:

- If a backup-archive client is installed on a client node, and the password is lost or forgotten, complete the following steps:

1. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwnreset=yes
```

where *node_name* specifies the client node and *new_password* specifies the password that you assign.

2. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to *generate* in the client options file.

- If an administrator is locked out because of password issues, complete the following steps:

1. To provide the administrator with access to the server, issue the UNLOCK ADMIN command. For instructions, see UNLOCK ADMIN (Unlock an administrator).
2. Set a new password by using the UPDATE ADMIN command:

```
update admin admin_name new_password forcepwnreset=yes
```

where *admin_name* specifies the name of the administrator and *new_password* specifies the password that you assign.

- If a client node is locked, complete the following steps:

1. Determine why the client node is locked and whether it must be unlocked. For example, if the client node is decommissioned, the client node is being removed from the production environment. You cannot reverse the decommission operation, and the client node remains locked. A client node also might be locked if the client data is the subject of a legal investigation.
2. If you must unlock a client node, use the UNLOCK NODE command. For instructions, see UNLOCK NODE (Unlock a client node).
3. Generate a new password by issuing the UPDATE NODE command:

```
update node node_name new_password forcepwnreset=yes
```

where *node_name* specifies the name of the node and *new_password* specifies the password that you assign.

4. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

Tip: The password is generated automatically if you previously set the passwordaccess option to *generate* in the client options file.

Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to the IBM Spectrum Protect™ server, but the workstation is no longer used, you can decommission the workstation.

About this task

When you start the decommission process, the server locks the client node to prevent it from accessing the server. Files that belong to the client node are gradually deleted, and then the client node is deleted. You can decommission the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes

Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a file space.

The simplest method for decommissioning a client node is to use the Operations Center. The decommission process runs in the background. If the client is configured to replicate client data, the Operations Center automatically removes the client from replication on the source and target replication servers before it decommissions the client.

Tip: Alternatively, you can decommission a client node by issuing the `DECOMMISSION NODE` or `DECOMMISSION VM` command. You might want to use this method in the following cases:

- To schedule the decommission process for the future or to run a series of commands by using a script, specify the decommission process to run in the background.
- To monitor the decommission process for debugging purposes, specify the decommission process to run in the foreground. If you run the process in the foreground, you must wait for the process to be completed before you continue with other tasks.

Procedure

Take one of the following actions:

- To decommission a client in the background by using the Operations Center, complete the following steps:
 1. On the Operations Center Overview page, click Clients and select the client.
 2. Click More > Decommission.
- To decommission a client node by using an administrative command, complete the following steps:
 1. Determine whether the client node is configured for node replication by issuing the `QUERY NODE` command. For example, if the client node is named AUSTIN, run the following command:

```
query node austin format=detailed
```

Review the Replication State output field.

2. If the client node is configured for replication, remove the client node from replication by issuing the `REMOVE REPLNODE` command. For example, if the client node is named AUSTIN, issue the following command:

```
remove replnode austin
```

3. Take one of the following actions:

- To decommission an application or system client node in the background, issue the `DECOMMISSION NODE` command. For example, if the client node is named AUSTIN, issue the following command:

```
decommission node austin
```

- To decommission an application or system client node in the foreground, issue the `DECOMMISSION NODE` command and specify the `wait=yes` parameter. For example, if the client node is named AUSTIN, issue the

following command:

```
decommission node austin wait=yes
```

- To decommission a virtual machine in the background, issue the DECOMMISSION VM command. For example, if the virtual machine is named AUSTIN, the file space is 7, and the file space name is specified by the file space ID, issue the following command:

```
decommission vm austin 7 nametype=fsid
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid
```

- To decommission a virtual machine in the foreground, issue the DECOMMISSION VM command and specify the `wait=yes` parameter. For example, issue the following command:

```
decommission vm austin 7 nametype=fsid wait=yes
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example:

```
decommission vm "austin 2" 7 nametype=fsid wait=yes
```

What to do next

Watch for error messages, which might be displayed in the user interface or in the command output, immediately after you run the process.

You can verify that the client node is decommissioned:

1. On the Operations Center Overview page, click Clients.
2. In the Clients table, in the At risk column, review the state:
 - A DECOMMISSIONED state specifies that the node is decommissioned.
 - A null value specifies that the node is not decommissioned.
 - A PENDING state specifies that the node is being decommissioned, or the decommission process failed.

Tip: If you want to determine the status of a pending decommission process, issue the following command:

```
query process
```

3. Review the command output:
 - If status is provided for the decommission process, the process is in progress. For example:

```
query process
```

Process Number	Process Description	Process Status
3	DECOMMISSION NODE	Number of backup objects deactivated for node NODE1: 8 objects deactivated.

- If no status is provided for the decommission process, and you did not receive an error message, the process is incomplete. A process can be incomplete if files that are associated with the node are not yet deactivated. After the files are deactivated, run the decommission process again.
- If no status is provided for the decommission process, and you receive an error message, the process failed. Run the decommission process again.

Related reference:

- [DECOMMISSION NODE \(Decommission a client node\)](#)
- [DECOMMISSION VM \(Decommission a virtual machine\)](#)
- [QUERY NODE \(Query nodes\)](#)
- [REMOVE REPLNODE \(Remove a client node from replication\)](#)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Spectrum Protect™ server. When you run the deactivation process, any backup data that was stored before the specified date and time is deactivated and will be deleted as it expires. In this way, you can free space on the server.

About this task

Some application clients always save data to the server as active backup data. Because active backup data is not managed by inventory expiration policies, the data is not deleted automatically, and uses server storage space indefinitely. To free the storage space that is used by obsolete data, you can deactivate the data.

When you run the deactivation process, all active backup data that was stored before the specified date becomes inactive. The data is deleted as it expires and cannot be restored. The deactivation feature applies only to application clients that protect Oracle databases.

Procedure

1. From the Operations Center Overview page, click Clients.
2. In the Clients table, select one or more clients and click More > Clean Up.
Command-line method: Deactivate data by using the DEACTIVATE DATA command.

Related reference:

[DEACTIVATE DATA \(Deactivate data for a client node\)](#)

Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Before you begin

1. Review the client/server compatibility requirements in technote 1053218. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted.
2. Verify system requirements for the client in IBM Spectrum Protect™ Supported Operating Systems.
3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with servers that are configured as library managers. See technote 1302789.

If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, complete the instructions that are listed in the following table.

Software	Link to instructions
IBM Spectrum Protect backup-archive client	<ul style="list-style-type: none">• Scheduling client updates
IBM Spectrum Protect Snapshot	<ul style="list-style-type: none">• Installing and upgrading IBM Spectrum Protect Snapshot for UNIX and Linux• Installing and upgrading IBM Spectrum Protect Snapshot for VMware• Installing and upgrading IBM Spectrum Protect Snapshot for Windows
IBM Spectrum Protect for Databases	<ul style="list-style-type: none">• Upgrading Data Protection for SQL Server• Data Protection for Oracle installation• Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Enterprise Resource Planning	<ul style="list-style-type: none">• Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for DB2®• Upgrading IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP for Oracle

Software	Link to instructions
IBM Spectrum Protect for Mail	<ul style="list-style-type: none"> • Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (V7.1.0) • Installation of Data Protection for IBM Domino on a Windows system (V7.1.0) • Installing, upgrading, and migrating IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
IBM Spectrum Protect for Virtual Environments	<ul style="list-style-type: none"> • Installing and upgrading Data Protection for VMware • Installing Data Protection for Microsoft Hyper-V

Managing the Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect™ environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

- Adding and removing spoke servers
In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.
- Starting and stopping the web server
The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.
- Restarting the initial configuration wizard
You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.
- Changing the hub server
You can use the Operations Center to remove the hub server of IBM Spectrum Protect, and configure another hub server.
- Restoring the configuration to the preconfiguration state
If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect servers are not defined as hub or spoke servers.

Adding and removing spoke servers

In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

About this task

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

- Adding a spoke server
After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.
- Removing a spoke server
You can remove a spoke server from the Operations Center.

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of the hub server.

Procedure

1. In the Operations Center menu bar, click Servers. The Servers page opens.

In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the DEFINE SERVER command, the server is not yet configured as a spoke server.

2. Complete one of the following steps:
 - o Click the server to highlight it, and in the table menu bar, click Monitor Spoke.
 - o If the server that you want to add is not shown in the table, and secure SSL/TLS communication is not required, click + Spoke in the table menu bar.
3. Provide the necessary information, and complete the steps in the spoke configuration wizard.
Tip: If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Removing a spoke server

You can remove a spoke server from the Operations Center.

About this task

You might need to remove a spoke server in the following situations, for example:

- You want to move the spoke server from one hub server to another hub server.
- You want to decommission the spoke server.

Procedure

To remove the spoke server from the group of servers that are managed by the hub server, complete the following steps:

1. From the IBM Spectrum Protect™ command line, issue the following command on the hub server:

```
QUERY MONITORSETTINGS
```

2. From the output of the command, copy the name that is in the Monitored Group field.
3. Issue the following command on the hub server, where *group_name* represents the name of the monitored group, and *member_name* represents the name of the spoke server:

```
DELETE GRPMEMBER group_name member_name
```

4. Optional: If you want to move the spoke server from one hub server to another hub server, do **not** complete this step. Otherwise, you can disable alerting and monitoring on the spoke server by issuing the following commands on the spoke server:

```
SET STATUSMONITOR OFF  
SET ALERTMONITOR OFF
```

5. Optional: If the spoke server definition is used for other purposes, such as enterprise configuration, command routing, storing virtual volumes, or library management, do **not** complete this step. Otherwise, you can delete the spoke server definition on the hub server by issuing the following command on the hub server:

```
DELETE SERVER spoke_server_name
```

Tip: If a server definition is deleted immediately after the server is removed from the monitored group, status information for the server can remain in the Operations Center indefinitely.

To avoid this issue, wait until the status collection interval passes before you delete the server definition. The status collection interval is shown on the Settings page of the Operations Center.

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might have to stop and start the web server, for example, to make configuration changes.

Procedure

1. Stop the web server.

- o **AIX** From the `/installation_dir/ui/utls` directory, where `installation_dir` represents the directory where the Operations Center is installed, issue the following command:


```
./stopserver.sh
```
 - o **Linux** Issue the following command:


```
service opscenter.rc stop
```
 - o **Windows** From the Services window, stop the IBM Spectrum Protect™ Operations Center service.
2. Start the web server.
- o **AIX** From the `/installation_dir/ui/utls` directory, where `installation_dir` represents the directory where the Operations Center is installed, issue the following command:


```
./startserver.sh
```
 - o **Linux** Issue the following commands:

Start the server:

```
service opscenter.rc start
```

Restart the server:

```
service opscenter.rc restart
```

Determine whether the server is running:

```
service opscenter.rc status
```
 - o **Windows** From the Services window, start the IBM Spectrum Protect Operations Center service.

Restarting the initial configuration wizard

You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.

Before you begin

To change the following settings, use the Settings page in the Operations Center rather than restarting the initial configuration wizard:

- The frequency at which status data is refreshed
- The duration that alerts remain active, inactive, or closed
- The conditions that indicate that clients are at risk

The Operations Center help includes more information about how to change these settings.

About this task

To restart the initial configuration wizard, you must delete a properties file that includes information about the hub server connection. However, any alerting, monitoring, at-risk, or multiserver settings that were configured for the hub server are not deleted. These settings are used as the default settings in the configuration wizard when the wizard restarts.

Procedure

1. Stop the Operations Center web server.
2. On the computer where the Operations Center is installed, go to the following directory, where `installation_dir` represents the directory in which the Operations Center is installed:
 - o **AIX** **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer`
 - o **Windows** `installation_dir\ui\Liberty\usr\servers\guiServer`

For example:

 - o **AIX** **Linux** `/opt/tivoli/tsm/ui/Liberty/usr/servers/guiServer`
 - o **Windows** `c:\Program Files\Tivoli\TSM\ui\Liberty\usr\servers\guiServer`
3. In the `guiServer` directory, delete the `serverConnection.properties` file.
4. Start the Operations Center web server.
5. Open the Operations Center.

6. Use the configuration wizard to reconfigure the Operations Center. Specify a new password for the monitoring administrator ID.
7. On any spoke servers that were previously connected to the hub server, update the password for the monitoring administrator ID by issuing the following command from the IBM Spectrum Protect™ command-line interface:

```
UPDATE ADMIN IBM-OC-hub_server_name new_password
```

Restriction: Do not change any other settings for this administrator ID. After you specify the initial password, this password is managed automatically by the Operations Center.

Changing the hub server

You can use the Operations Center to remove the hub server of IBM Spectrum Protect™, and configure another hub server.

Procedure

1. Restart the initial configuration wizard of the Operations Center. As part of this procedure, you delete the existing hub server connection.
2. Use the wizard to configure the Operations Center to connect to the new hub server.

Related tasks:

Restarting the initial configuration wizard

Restoring the configuration to the preconfiguration state

If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect™ servers are not defined as hub or spoke servers.

Procedure

To restore the configuration, complete the following steps:

1. Stop the Operations Center web server.
2. Unconfigure the hub server by completing the following steps:
 - a. On the hub server, issue the following commands:

```
SET MONITORINGADMIN ""
SET MONITOREDSEVERGROUP ""
SET STATUSMONITOR OFF
SET ALERTMONITOR OFF
REMOVE ADMIN IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. Reset the password for the hub server by issuing the following command on the hub server:

```
SET SERVERPASSWORD ""
```

Attention: Do not complete this step if the hub server is configured with other servers for other purposes, such as library sharing, exporting and importing of data, or node replication.

3. Unconfigure any spoke servers by completing the following steps:
 - a. On the hub server, to determine whether any spoke servers remain as members of the server group, issue the following command:

```
QUERY SERVERGROUP IBM-OC-hub_server_name
```

Tip: *IBM-OC-hub_server_name* represents the name of the monitored server group that was automatically created when you configured the first spoke server. This server group name is also the same as the monitoring administrator ID that was automatically created when you initially configured the hub server.

- b. On the hub server, to delete spoke servers from the server group, issue the following command for each spoke server:

```
DELETE GRPMEMBER IBM-OC-hub_server_name spoke_server_name
```

- c. After all spoke servers are deleted from the server group, issue the following commands on the hub server:

```
DELETE SERVERGROUP IBM-OC-hub_server_name
SET MONITOREDSEVERGROUP ""
```

d. On each spoke server, issue the following commands:

```
REMOVE ADMIN IBM-OC-hub_server_name
SETOPT PUSHSTATUS NO
SET ALERTMONITOR OFF
SET STATUSMONITOR OFF
```

e. On each spoke server, delete the definition of the hub server by issuing the following command:

```
DELETE SERVER hub_server_name
```

Attention: Do not complete this step if the definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

f. On the hub server, delete the definition of each spoke server by issuing the following command:

```
DELETE SERVER spoke_server_name
```

Attention: Do not complete this step if the server definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

4. Restore the default settings on each server by issuing the following commands:

```
SET STATUSREFRESHINTERVAL 5
SET ALERTUPDATEINTERVAL 10
SET ALERTACTIVEDURATION 480
SET ALERTINACTIVEDURATION 480
SET ALERTCLOSEDDURATION 60
SET STATUSATRISKINTERVAL TYPE=AP INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=VM INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=SY INTERVAL=24
SET STATUSSKIPASFAILURE YES TYPE=ALL
```



5. Restart the initial configuration wizard of the Operations Center.

Related tasks:

Restarting the initial configuration wizard
Starting and stopping the web server

Configuring virtual tape libraries

A virtual tape library (VTL) does not use physical tape media. When you implement VTL storage, you can exceed the capacity of a physical tape library. The ability to define many volumes and drives can provide greater flexibility for the storage environment.

- Considerations for using virtual tape libraries
There are some considerations for defining a library as a virtual tape library (VTL), including enhancements for performance and setup of your hardware.
- Adding a virtual tape library to your environment
Define a virtual tape library (VTL) to take advantage of mount performance and scalability advantages.
- Defining all drives and paths for a single library
Use the `PERFORM LIBACTION` command to set up a single SCSI or virtual tape library (VTL) in one step.
-  Example: Configure a SCSI or virtual tape library with a single drive device type
Configure a VTL or SCSI library that contains two LTO tape drives.
-  Example: Configure a SCSI or virtual tape library with multiple drive device types
You can configure a library with multiple drive device types, for example, a StorageTek L40 library that contains one DLT drive and one LTO Ultrium drive.

Considerations for using virtual tape libraries

There are some considerations for defining a library as a virtual tape library (VTL), including enhancements for performance and setup of your hardware.

About this task

Defining a VTL to the IBM Spectrum Protect™ server can help improve performance because the server handles mount point processing for VTLs differently than real tape libraries. The physical limitations for real tape hardware are not applicable to a VTL, affording options for better scalability.

You can use a VTL for any virtual tape library when the following conditions are true:

- There is no mixed media involved in the VTL. Only one type and generation of drive and media is emulated in the library.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

If either of these conditions are not met, any mount performance advantage from defining a VTL library to the IBM Spectrum Protect server can be reduced or negated.

VTLs are compatible with earlier versions of both library clients and storage agents. The library client or storage agent is not affected by the type of library that is used for storage. If mixed media and path conditions are true for a SCSI library, it can be defined or updated as LIBTYPE=VTL.

- **Storage capacity for virtual tape libraries**
Because virtual tape libraries (VTLs) do not have the physical limitations that real tape hardware does, their capacity for storage is more flexible.
- **Drive configuration for virtual tape libraries**
Drive configuration in a virtual tape library (VTL) is variable, depending on the needs of your environment.

Storage capacity for virtual tape libraries

Because virtual tape libraries (VTLs) do not have the physical limitations that real tape hardware does, their capacity for storage is more flexible.

The concept of storage capacity in a virtual tape library is different from capacity in physical tape hardware. In a physical tape library, each volume has a defined capacity, and the library's capacity is defined in terms of the total number of volumes in the library. The capacity of a VTL, alternatively, is defined in terms of total available disk space. You can increase or decrease the number and size of volumes on disk.

This variability affects what it means to run out of space in a VTL. For example, a volume in a VTL can run out of space before reaching its assigned capacity if the total underlying disk runs out of space. In this situation, the server can receive an end-of-volume message without any warning, resulting in backup failures.

When out-of-space errors and backup failures occur, disk space is usually still available in the VTL. It is hidden in volumes that are not in use. For example, volumes that are logically deleted or returned to scratch status in the IBM Spectrum Protect™ server are deleted only in the server database. The VTL is not notified, and the VTL maintains the full size of the volume as allocated in its capacity considerations.

To help prevent out-of-space errors, ensure that any SCSI library that you update to LIBTYPE=VTL is updated with the RELABELSCRATCH parameter set to YES. The RELABELSCRATCH option enables the server to overwrite the label for any volume that is deleted and to return the volume to scratch status in the library. The RELABELSCRATCH parameter defaults to YES for any library defined as a VTL.

Related reference:

UPDATE LIBRARY (Update a library)

Drive configuration for virtual tape libraries

Drive configuration in a virtual tape library (VTL) is variable, depending on the needs of your environment.

Most VTL environments use as many drives as possible to maximize the number of concurrent tape operations. A single tape mount in a VTL environment is typically faster than a physical tape mount. However, using many drives increases the amount of time that the IBM Spectrum Protect™ server requires when a mount is requested. The selection process takes longer as the number of drives that are defined in a single library object in the server increases. Virtual tape mounts can take as long or longer than physical tape mounts depending on the number of drives in the VTL.

For best results when you create drives, check with your VTL vendor about device-specific recommendations. If more than 300-500 drives for each VTL are required, you can logically partition the VTL into multiple libraries and assign drives to each library. Operating system and SAN hardware configurations could impose limitations on the number of devices that can be utilized within the VTL library.

Adding a virtual tape library to your environment

Define a virtual tape library (VTL) to take advantage of mount performance and scalability advantages.

About this task

VTLs are identified by using the DEFINE LIBRARY command and specifying the LIBTYPE=VTL parameter. Because a VTL library functionally interacts with the server in the same way that a SCSI library does, you can use the UPDATE LIBRARY command to change the library type of a SCSI library that is already defined. You do not have to redefine the library.

Procedure

- Add a new VTL library. Define the library as a VTL to the server, as shown in the following example:

```
define library chester libtype=vtl
```

This sets up the new VTL library and enables the RELABELSCRATCH option to relabel volumes that have been deleted and returned to scratch status.

- Update a SCSI library to a VTL. If you have a SCSI library and you want to change it to a VTL, use the UPDATE LIBRARY command to change the library type:

```
update library calzone libtype=vtl
```

You can issue this command only if the library that is being updated is defined with the LIBTYPE=SCSI parameter.

Related reference:

DEFINE LIBRARY (Define a library)

UPDATE LIBRARY (Update a library)

Defining all drives and paths for a single library

Use the PERFORM LIBACTION command to set up a single SCSI or virtual tape library (VTL) in one step.

About this task

If you are setting up or modifying your hardware environment and must create or change large numbers of drive definitions, the PERFORM LIBACTION command can make this task much simpler. You can define a new library and then define all drives and paths to the drives. Or, if you have an existing library that you want to delete, you can delete all existing drives and their paths in one step.

The PREVIEW parameter allows you to view the output of commands before they are processed to verify the action that you want to perform. If you are defining a library, a path to the library must already be defined if you want to specify the PREVIEW parameter. You cannot use the PREVIEW and DEVICE parameters together.

The PERFORM LIBACTION command can be used only for SCSI and VTL libraries. If you are defining drives and paths for a library, the SANDISCOVERY option must be supported and enabled. The tape library must be able to return the drive serial number address association.

Procedure

To set up a VTL library named ODIN, complete these steps:

1. Define the library.

```
define library odin libtype=vtl
```

2. Define two drives and their paths for your new library, ODIN.

AIX

```
perform libaction odin action=define device=/dev/lb3 prefix=dr
```

The server then issues the following commands:


```
define path tsmserver odin srct=server destt=library device=/dev/
lb3 define drive odin dr0
define path tsmserver dr0 srct=server destt=drive library=odin
device=/dev/mt1 define drive odin dr1
define path tsmserver dr1 srct=server destt=drive library=odin
device=/dev/mt2
```

Linux

```
perform libaction odin action=define device=/dev/tmscsi/lb3 prefix=dr
```

The server then issues the following commands:

```
define path tsmserver odin srct=server destt=library device=/dev/tmscsi/lb3
define drive odin dr0
define path tsmserver dr0 srct=server destt=drive library=odin
device=/dev/tmscsi/mt1 define drive odin dr1
define path tsmserver dr1 srct=server destt=drive library=odin
device=/dev/tmscsi/mt2
```

Windows

```
perform libaction odin action=define device=lb0.0.0.2 prefix=dr
```

The server then issues the following commands:

```
define path tsmserver odin srct=server destt=library device=lb0.0.0.2
define drive odin dr0
define path tsmserver dr0 srct=server destt=drive library=odin
device=mt0.1.0.2 define drive odin dr1
define path tsmserver dr1 srct=server destt=drive library=odin
device=mt0.2.0.2
```

Related reference:

DEFINE LIBRARY (Define a library)

DEFINE PATH (Define a path when the destination is a drive)

PERFORM LIBACTION (Define or delete all drives and paths for a library)

Example: Configure a SCSI or virtual tape library with a single drive device type

Configure a VTL or SCSI library that contains two LTO tape drives.

About this task

This procedure is an example of configuring an automated SCSI library that contains two drives to the server system. The library is not shared with other IBM Spectrum Protect™ servers or with storage agents and is typically attached to the server system by using SCSI cables.

In this configuration, both drives in the library are the same device type. Define one device class. The procedure is the same for both SCSI libraries and VTLs, except for the step to define the library. For SCSI libraries, define the library with libtype=scsi. For VTL libraries, define the library with libtype=vtl.

Procedure

1. Define a SCSI library that is named AUTODTLIB.

```
define library autoltolib libtype=scsi
```

If the library has a bar code reader and you would like to automatically label tapes before they are checked in, you can set the AUTOLABEL parameter to YES. For example:

```
define library autoltolib libtype=scsi autolabel=yes
```

2. Define a path from the server to the library.

AIX

```
define path server1 autoltolib srctype=server desttype=library
device=/dev/lb3
```

Linux

```
define path server1 autoltolib srctype=server desttype=library
device=/dev/tmscsi/lb3
```

Windows

```
define path server1 autoltolib srctype=server desttype=library
device=lb0.0.0.3
```

3. Define the drives in the library. Both drives belong to the AUTODTLIB library.

```
define drive autoltolib drive01
define drive autoltolib drive02
```

Tip: You can use the `PERFORM LIBACTION` command to define drives and paths for a library in one step.

4. Define a path from the server to each drive.

AIX

```
define path server1 drive01 srctype=server desttype=drive
library=autoltolib device=/dev/mt4
define path server1 drive02 srctype=server desttype=drive
library=autoltolib device=/dev/mt5
```

Linux

```
define path server1 drive01 srctype=server desttype=drive
library=autoltolib device=/dev/tmscsi/mt4
define path server1 drive02 srctype=server desttype=drive
library=autoltolib device=/dev/tmscsi/mt5
```

Windows

```
define path server1 drive01 srctype=server desttype=drive
library=autoltolib device=mt0.0.0.4
define path server1 drive02 srctype=server desttype=drive
library=autoltolib device=mt0.0.0.5
```

If you did not include the element address when you defined the drive, the server now queries the library to obtain the default element address for the drive.

5. Define a device class that is named `AUTODLT_CLASS` for the two drives in the `AUTODTLIB` library.

```
define devclass autolto_class library=autodltlib devtype=lto
```

6. Define a storage pool that is named `AUTOLTO_POOL` associated with the device class named `AUTOLTO_CLASS`.

```
define stgpool autolto_pool autolto_class maxscratch=20
```

7. Label and check in library volumes.

```
label libvolume autoltolib search=yes labelsource=barcode checkin=scratch
```

8. Verify your definitions by issuing the following commands:

```
query library
query drive
query path
query devclass
query stgpool
query libvolume
```

Related reference:

`DEFINE DEVCLASS` (Define a device class)

`DEFINE LIBRARY` (Define a library)

`DEFINE PATH` (Define a path when the destination is a drive)

Example: Configure a SCSI or virtual tape library with multiple drive device types

You can configure a library with multiple drive device types, for example, a StorageTek L40 library that contains one DLT drive and one LTO Ultrium drive.

About this task

This procedure is an example of configuring an automated SCSI library that contains two drives to the server system. The library is not shared with other IBM Spectrum Protect™ servers or with storage agents and is typically attached to the server system by SCSI cables.

In this configuration, the drives are different device types. Define a device class for each drive device type. Drives with different device types are supported in a single library if you define a device class for each type of drive. If you are configuring this way, you must include the specific format for the drive's device type by using the FORMAT parameter with a value other than DRIVE.

The procedure is the same for both SCSI libraries and VTLs, except for the step to define the library. For SCSI libraries, define the library with `libtype=scsi`. For VTL libraries, define the library with `libtype=vtl`.

Procedure

1. Define a SCSI library named MIXEDLIB.

```
define library mixedlib libtype=scsi
```

2. Define a path from the server to the library.

AIX

```
define path server1 mixedlib srctype=server desttype=library  
device=/dev/lb3
```

Linux

```
define path server1 mixedlib srctype=server desttype=library  
device=/dev/tmscsi/lb3
```

Windows

```
define path server1 mixedlib srctype=server desttype=library  
device=lb0.0.0.3
```

3. Define the drives in the library. Both drives belong to the MIXEDLIB library.

```
define drive mixedlib dlt1  
define drive mixedlib lto1
```

4. Define a path from the server to each drive. The DEVICE parameter specifies the device driver's name for the drive, which is the device special file name.

AIX

```
define path server1 dlt1 srctype=server desttype=drive  
library=mixedlib device=/dev/mt4  
define path server1 lto1 srctype=server desttype=drive  
library=mixedlib device=/dev/mt5
```

Linux

```
define path server1 dlt1 srctype=server desttype=drive  
library=mixedlib device=/dev/tmscsi/mt4  
define path server1 lto1 srctype=server desttype=drive  
library=mixedlib device=/dev/tmscsi/mt5
```

Windows

```
define path server1 drive01 srctype=server desttype=drive  
library=autoltolib device=mt0.0.0.4  
define path server1 drive02 srctype=server desttype=drive  
library=autoltolib device=mt0.0.0.5
```

If you did not include the element address when you defined the drive, the server now queries the library to obtain the element address for the drive.

5. Define device classes.

Important: Do not use the DRIVE format, which is the default. Because the drives are different types, the server uses the format specification to select a drive. The results of using the DRIVE format in a mixed media library are unpredictable.

```
define devclass dlt_class library=mixedlib devtype=dlt format=dlt40  
define devclass lto_class library=mixedlib devtype=lto format=ultriumc
```

6. Define storage pools that are associated with the device classes.

```
define stgpool lto_pool lto_class maxscratch=20
define stgpool dlt_pool dlt_class maxscratch=20
```

7. Label and check in library volumes.

```
label libvolume mixedlib search=yes labelsource=barcode checkin=scratch
```

8. Verify your definitions by issuing the following commands:

```
query library
query drive
query path
query devclass
query stgpool
query libvolume
```

Protecting NAS file servers

You can configure and manage a backup environment that protects a network-attached storage (NAS) file server.

You can use the IBM Spectrum Protect™ server, the IBM Spectrum Protect backup-archive client, or IBM Spectrum Protect Snapshot to back up and restore a NAS file server as described in the following table.

Product	Description
IBM Spectrum Protect server	<p>To back up and restore NAS file server data by using the IBM Spectrum Protect server, you must have IBM Spectrum Protect Extended Edition installed.</p> <p>You can configure the IBM Spectrum Protect server to use the network data management protocol (NDMP) to back up and restore data as described in the following topics in this section.</p> <p>To protect large NetApp file systems, you can alternatively configure IBM Spectrum Protect to use NetApp SnapMirror to Tape (also known as SMTape). SnapMirror to Tape uses a block-level copy of data for backup, which is faster than a traditional NDMP full backup and can be used when NDMP full backups are impractical.</p> <p>For information about using SnapMirror to Tape to back up and restore data, see Backup and restore operations by using the NetApp SnapMirror to Tape feature.</p>
IBM Spectrum Protect backup-archive client	<p>You can configure the backup-archive client to back up and restore file server data by using the Network File System (NFS) or Common Internet File System (CIFS) protocol.</p> <p>For information about using the backup-archive client to back up and restore data, see Back up and restore data with backup-archive clients.</p>
IBM Spectrum Protect Snapshot	<p>You can use IBM Spectrum Protect Snapshot to back up and restore file server data by using the advanced snapshot technologies of storage systems.</p> <p>For information about using IBM Spectrum Protect Snapshot to back up and restore data, see IBM Spectrum Protect Snapshot for UNIX and Linux overview or IBM Spectrum Protect Snapshot for VMware overview.</p>

- NDMP requirements
To use NDMP for operations with NAS file servers, you must have IBM Spectrum Protect Extended Edition installed and your file server environment must meet certain requirements.
- NDMP operations management
There are several administrator activities for NDMP operations.
- Configuring IBM Spectrum Protect for NDMP operations
You can configure IBM Spectrum Protect to back up and recover data on NAS file servers by using NDMP. The configuration procedure differs depending on whether you plan to back up data from a nonclustered or clustered NAS file server.
- Backing up and restoring NAS file servers using NDMP
After you configure IBM Spectrum Protect for NDMP operations, you are ready to begin using NDMP.
- File-level backup and restore for NDMP operations
When you back up data by using NDMP, you can specify that the IBM Spectrum Protect server collects and stores file-level information in a table of contents (TOC).
- Directory-level backup and restore operations
If you have a large NAS file system, initiating a backup at a directory level reduces backup and restore times and provides more flexibility in configuring NAS backups. By defining virtual file spaces, a file system backup can be partitioned among

several NDMP backup operations and multiple tape drives. You can also use different backup schedules to back up subtrees of a file system.

- Backup and restore operations by using the NetApp SnapMirror to Tape feature
You can back up large NetApp file systems by using the NetApp SnapMirror to Tape feature (also known as SMTape). Using a block-level copy of data for backup, the SnapMirror to Tape method is faster than a traditional NDMP full backup and can be used when NDMP full backups are impractical.
- NDMP backup operations using Celerra file server-integrated checkpoints
When the IBM Spectrum Protect server initiates an NDMP backup operation on a Celerra data mover, the backup of a large file system might take several hours to complete. Without Celerra integrated checkpoints, any changes that occur on the file system are written to the backup image.
- Replicating NAS nodes
You can replicate a NAS node that uses NDMP for backup operations. Before you configure the replication operation, review the restrictions that apply.

NDMP requirements

To use NDMP for operations with NAS file servers, you must have IBM Spectrum Protect™ Extended Edition installed and your file server environment must meet certain requirements.

NAS file server

The operating system on the file server must be supported by IBM Spectrum Protect. For information about the NAS file servers that are supported, see technote 1054144.

The combination of file server model and operating system must be supported by the NAS file server. For more specifics, see the product information for the NAS file server.

Tape libraries

This requirement is necessary only for a backup to a locally attached NAS device. The IBM Spectrum Protect server supports the following types of libraries for operations that use NDMP:

SCSI

A SCSI library can be attached directly to the IBM Spectrum Protect server or to the NAS file server. When the library is attached directly to the IBM Spectrum Protect server, that server controls the library operations by passing the SCSI commands directly to the library. When the library is attached directly to the NAS file server, the IBM Spectrum Protect server controls the library by passing SCSI commands to the library through the NAS file server.

ACSLs

An automated cartridge system library software (ACSLs) library can be directly connected only to the IBM Spectrum Protect server. The IBM Spectrum Protect server controls the library by passing the library request through TCP/IP to the library control server.

Restriction: The IBM Spectrum Protect server does not include External Library support for the ACSLS library when the library is used for NDMP operations.

VTL

A virtual tape library (VTL) can be attached directly either to the IBM Spectrum Protect server or to the NAS file server. A virtual tape library is essentially the same as a SCSI library, but is enhanced for virtual tape library characteristics and allows for better mount performance.

If you are defining a VTL, your environment must not include mixed media. Paths must be defined between all drives in the library and all defined servers, including storage agents, that use the library. If these conditions are not met, the overall performance can degrade to the same levels as the SCSI library type, especially during times of high stress.

349X

A 349X library can be directly connected only to the IBM Spectrum Protect server. The IBM Spectrum Protect server controls the library by passing the library request through TCP/IP to the library manager.

Library sharing: The IBM Spectrum Protect server that runs NDMP operations can be a library manager for either an ACSLS, SCSI, VTL, or 349X library, but cannot be a library client. The IBM Spectrum Protect server can also be a library client in a configuration where the NAS file server sends data to the server by using TCP/IP rather than to a tape library attached to the file server. If the IBM Spectrum Protect server that runs NDMP operations is a library manager, that server must control the library directly and not by passing commands through the NAS file server.

Tape drives

A tape drive is necessary only for backup to a locally attached NAS device. The NAS file server must be able to access the drives. A NAS device is not supported in a mixed device library. The drives must be supported for tape backup operations by

the NAS file server and its operating system. For complete NDMP device support, refer to the NAS file server product documentation.

Drive sharing: The tape drives can be shared by the IBM Spectrum Protect server and one or more NAS file servers. Also, when a SCSI, VTL, or a 349X library is connected to the server and not to the NAS file server, the drives can be shared by one or more NAS file servers. The drives can also be shared by one or more IBM Spectrum Protect library clients and storage agents.

Drive reservations: When tape drives are attached to NAS devices and the RESETDRIVES=YES parameter for the DEFINE LIBRARY command is specified, the following limitations apply:

- If a tape drive is shared by an IBM Spectrum Protect server and a NAS device, drive reservation preemption is supported when the NAS device supports persistent reserve and it is enabled. For more information about setting persistent reserve, see the documentation for your NAS device.
- If a tape drive is attached only to a NAS device and not shared with an IBM Spectrum Protect server, drive reservation preemption is not supported. If you enable persistent reserve on the NAS device for these drives and a reservation is set by the NAS device but never cleared, you must use another method to clear the reservation.

Verify the compatibility of specific combinations of a NAS file server, tape devices, and SAN-attached devices with the hardware manufacturers.

Tip: IBM Spectrum Protect supports NDMP Version 4 for all NDMP operations. IBM Spectrum Protect continues to support all NDMP backup and restore operations with a NAS device that runs NDMP version 3. The IBM Spectrum Protect server negotiates the highest protocol level (either Version 3 or Version 4) with the NDMP server when it establishes an NDMP connection. If you experience any issues with Version 4, you might want to try Version 3.

- Interfaces for NDMP operations
You can use several interfaces to run NDMP operations. You can schedule an NDMP operation by using the BACKUP NODE or RESTORE NODE command and by creating a schedule to process the command.
- Data formats for NDMP backup operations
Data that is backed up using NDMP is not in the same format as the data that is used for typical IBM Spectrum Protect backup operations. The NAS file server controls the format of the backup data.
- Storage pool types for NDMP operations
Before you configure IBM Spectrum Protect for network data management protocol (NDMP) operations, review the storage pool types that are supported. Different storage pool types are supported for NDMP operations, depending on the brand of file server that you use.

Interfaces for NDMP operations

You can use several interfaces to run NDMP operations. You can schedule an NDMP operation by using the BACKUP NODE or RESTORE NODE command and by creating a schedule to process the command.

Client interfaces:

- IBM Spectrum Protect™ backup-archive command-line client (on a Windows, 64-bit AIX®, or 64-bit Oracle Solaris system)
- IBM Spectrum Protect web client interface, available with the backup-archive client, at Version 8.1.1 or earlier

Restriction: If you installed the backup-archive client V8.1.1 or earlier, you can use the web client interface for file-level restore operations. If you installed the backup-archive client V8.1.2 or later, you cannot use the web client interface for file-level restore operations.

Server interfaces:

- Server console
 - Command line on the administrative client
- Tip: All examples for NDMP operations use server commands.

The web client interface V8.1.1 or earlier displays the file systems of the NAS file server in a graphical view. The client function is not required, but you can use the client interfaces for NDMP operations. For file-level restore operations, the preferred method is to use the web client interface V8.1.1 or earlier. For more information about file-level restore operations, see File-level backup and restore for NDMP operations.

IBM Spectrum Protect prompts you for an administrator ID and password when you complete NDMP functions by using either of the client interfaces. For more information about installing and activating client interfaces, see Installing the IBM Spectrum Protect backup-archive clients.

To use the IBM Spectrum Protect backup-archive client or web client for NAS operations, the file system names on the NAS device must have a forward slash (/) as the first character. This restriction does not affect NAS operations that are initiated from the IBM

Data formats for NDMP backup operations

Data that is backed up using NDMP is not in the same format as the data that is used for typical IBM Spectrum Protect™ backup operations. The NAS file server controls the format of the backup data.

Data that is backed up to a library that is directly attached to the file server must be directed to a storage pool with the proper data format. When you define a storage pool for NDMP operations, you specify one of the following data formats:

- NETAPPDUMP if the NAS file server is a NetApp or an IBM® System Storage® N Series device.
- CELERRADUMP if the NAS file server is an EMC Celerra device.
- NDMPDUMP for all other devices.

Data that is backed up over the network to the local IBM Spectrum Protect hierarchy can be directed to any random-access or sequential-access primary storage pool. However, the format of the data does not change.

Storage pool types for NDMP operations

Before you configure IBM Spectrum Protect™ for network data management protocol (NDMP) operations, review the storage pool types that are supported. Different storage pool types are supported for NDMP operations, depending on the brand of file server that you use.

Backup operations

The following storage pool types can be used for backup operations.

Brand of file server	Can directory-container storage pools be used as the destination?	Can cloud-container storage pools be used as the destination?	Can non-container, non-deduplicated storage pools be used as the destination?	Can non-container, deduplicated storage pools of type FILE be used as the destination?
NetApp without the SnapMirror to Tape feature	Yes	Yes	Yes	Yes
NetApp with the SnapMirror to Tape feature	No	No	Yes	No
Other brands	No	No	Yes	No

Replication operations: limitations on source storage pools

The following storage pool types can be used on a source server for replication operations.

Brand of file server	Can directory-container storage pools be used on the source replication server?	Can cloud-container storage pools be used on the source replication server?	Can non-container, non-deduplicated storage pools be used on the source replication server?	Can non-container, deduplicated storage pools of type FILE be used on the source replication server?
NetApp without the SnapMirror to Tape feature	No	No	Yes	No
NetApp with the SnapMirror to Tape feature	No	No	Yes	No
Other brands	No	No	Yes	No

Replication operations: limitations on target storage pools

The following storage pool types can be used on a target replication server.

Brand of file server	Can directory-container storage pools be used on the target replication server?	Can cloud-container storage pools be used on the target replication server?	Can non-container, non-deduplicated storage pools be used on the target replication server?	Can non-container, deduplicated storage pools of type FILE be used on the target replication server?
NetApp without the SnapMirror to Tape feature	Yes	Yes	Yes	Yes
NetApp with the SnapMirror to Tape feature	No	No	Yes	No
Other brands	No	No	Yes	No

Protection operations to a remote storage pool

The following storage pool types can be used on a target replication server to protect data in directory-container storage pools by using the PROTECT STGPPOOL command.

Brand of file server	Can directory-container storage pools be used as the destination?	Can cloud-container storage pools be used as the destination?	Can non-container, non-deduplicated storage pools be used as the destination?	Can non-container, deduplicated storage pools of type FILE be used as the destination?
NetApp without the SnapMirror to Tape feature	Yes	N/A	N/A	N/A
NetApp with the SnapMirror to Tape feature	No	N/A	N/A	N/A
Other brands	No	N/A	N/A	N/A

Protection operations to tape on the same server

The following storage pool types can be used when you run the PROTECT STGPPOOL command to protect a directory-container storage pool to tape on the same server.

Brand of file server	Can directory-container storage pools be used as the destination?	Can cloud-container storage pools be used as the destination?	Can non-container, non-deduplicated storage pools be used as the destination?	Can non-container, deduplicated storage pools of type FILE be used as the destination?
NetApp without the SnapMirror to Tape feature	Yes	N/A	N/A	N/A
NetApp with the SnapMirror to Tape feature	No	N/A	N/A	N/A
Other brands	No	N/A	N/A	N/A

Storage pool conversion

If NDMP data exists in a storage pool that is converted to a directory-container storage pool or a cloud-container storage pool, the NDMP data remains in the original storage pool and is not converted.

Cloud tiering

If NDMP data exists in a storage pool that is tiered to cloud object storage, the NDMP data remains in the original storage pool and is not tiered.

Restrictions also apply when NETAPPDUMP, CELERRADUMP, or NDMPDUMP is designated as the storage-pool type. For more information, see Storage pool management for NDMP operations.

NDMP operations management

There are several administrator activities for NDMP operations.

- Managing NAS file server nodes
You can query, update, rename, and remove NAS file server nodes.
- Managing data movers that are used in NDMP operations
You can query, update, and delete the data movers that you define for NAS file servers.
- Dedicating an IBM Spectrum Protect drive to NDMP operations
If you are already using a drive for IBM Spectrum Protect™ operations, you can dedicate that drive to NDMP operations.
- Storage pool management for NDMP operations
When NETAPPDUMP, CELERRADUMP, or NDMPDUMP is designated as the storage-pool type, managing the storage pools that are produced by NDMP operations is different from managing storage pools that contain media for traditional IBM Spectrum Protect backups.
- Managing tables of contents
You can use several commands to manage different aspects of your data contents.
- Preventing long-running, inactive NDMP connections from closing
To prevent firewalls from closing NDMP connections that are long-running but inactive, you can enable Transmission Control Protocol (TCP) keepalive on the NDMP control connections.

Managing NAS file server nodes

You can query, update, rename, and remove NAS file server nodes.

Procedure

Use one of the following commands to manage NAS file server nodes:

Command	Procedure
QUERY NODE	To query a node, issue the QUERY NODE command with the appropriate parameters. For example, if you want to query the NAS node NASNODE1, issue the following command: <code>query node nasnode1 type=nas</code>
UPDATE NODE	To update a node, issue the UPDATE NODE command with the appropriate parameters. For example, if you created a new policy domain that is named NASDOMAIN for NAS nodes and you want to update the node NASNODE1 to include the node in the new domain, issue the following command: <code>update node nasnode1 domain=nasdomain</code>

Command	Procedure
RENAME NODE	<p>To rename a NAS node, you must also rename the corresponding NAS data mover; both must have the same name.</p> <p>For example, to rename NASNODE1 to NAS1, complete the following steps:</p> <ol style="list-style-type: none"> 1. Delete all paths between the data mover NASNODE1 and libraries and between the data mover NASNODE1 and drives. 2. Delete the data mover that is defined for the NAS node. 3. To rename NASNODE1 to NAS1, issue the following command: <pre>rename node nasnode1 nas1</pre> 4. Define the data mover by using the new node name. In this example, you must define a new data mover that is named NAS1 with the same parameters that were used to define NASNODE1. Important: When you define a new data mover for a node that you renamed, ensure that the data mover name matches the new node name. Also, ensure that the new data mover parameters are duplicates of the original data mover parameters. Any mismatch between a node name and a data mover name or between new data mover parameters and original data mover parameters can prevent you from establishing a session with the NAS file server. 5. For SCSI or 349X libraries, define a path between the NAS data mover and a library only if the tape library is physically connected directly to the NAS file server. 6. Define paths between the NAS data mover and any drives that are used for NDMP operations.
REMOVE NODE	<p>To remove a node, complete the following steps:</p> <ol style="list-style-type: none"> 1. Delete any virtual file space definitions for the node. 2. Delete all paths between the data mover and libraries and between the data mover and drives. 3. Delete the node. For example, if you want to remove a node that is named NAS1, issue the following command: <pre>remove node nas1</pre>

Related reference:

- QUERY NODE (Query nodes)
- UPDATE NODE (Update node attributes)
- RENAME NODE (Rename a node)
- REMOVE NODE (Delete a node or an associated machine node)

Managing data movers that are used in NDMP operations

You can query, update, and delete the data movers that you define for NAS file servers.

Procedure

Use one of the following commands to manage data movers:

Command	Procedure

Command	Procedure
QUERY DATAMOVER	To query a data mover, issue the QUERY DATAMOVER command with the appropriate parameters. For example, if you want to query the data mover NASNODE1, issue the following command: <code>query datamover nasnode1</code>
UPDATE DATAMOVER	To update a data mover, issue the UPDATE DATAMOVER command with the appropriate parameters. For example, if you shut down a NAS file server for maintenance and you want to take the data mover offline, issue the following command: <code>update datamover nasnode1 online=no</code>
DELETE DATAMOVER	To delete a data mover, issue the DELETE DATAMOVER command. For example, if you want to delete the data mover NASNODE1, issue the following command: <code>delete datamover nasnode1</code> Restriction: If the data mover has a path to a library, and you delete the data mover or take the data mover offline, you disable access to the library.

Related reference:

QUERY DATAMOVER (Display data mover definitions)

UPDATE DATAMOVER (Update a data mover)

DELETE DATAMOVER (Delete a data mover)

Dedicating an IBM Spectrum Protect drive to NDMP operations

If you are already using a drive for IBM Spectrum Protect™ operations, you can dedicate that drive to NDMP operations.

Procedure

Remove IBM Spectrum Protect server access by deleting the path definition. For example, if the server name is SERVER1 and the drive is NASDRIVE1, issue the following command:

```
delete path server1 nasdrive1 srctype=server desttype=drive library=naslib
```

Storage pool management for NDMP operations

When NETAPPDUMP, CELERRADUMP, or NDMPDUMP is designated as the storage-pool type, managing the storage pools that are produced by NDMP operations is different from managing storage pools that contain media for traditional IBM Spectrum Protect™ backups.

The following guidelines and restrictions apply to storage pools of the NETAPPDUMP, CELERRADUMP, and NDMPDUMP type that are produced by NDMP operations:

- You can query and update storage pools, but you cannot update the DATAFORMAT parameter.
- You cannot designate a CENTERA, directory-container, or cloud-container storage pool as a target pool of NDMP operations.
- Maintaining separate storage pools for data from different NAS vendors is the preferred practice, even when the data format for both is NDMPDUMP.
- The following DEFINE STGPOOL and UPDATE STGPOOL command parameters are ignored because storage pool hierarchies, reclamation, and migration are not supported for these storage pools:
 - MAXSIZE
 - NEXTSTGPOOL
 - LOWMIG
 - HIGHMIG
 - MIGDELAY
 - MIGCONTINUE
 - RECLAIMSTGPOOL

- o OVFLOCATION

Important: Ensure that you do not accidentally use storage pools that were defined for NDMP operations in traditional IBM Spectrum Protect operations. Be especially careful when you assign the storage pool name as the value for the DESTINATION parameter of the DEFINE COPYGROUP command. Unless the destination is a storage pool with the appropriate data format, the backup fails.

Managing tables of contents

You can use several commands to manage different aspects of your data contents.

About this task

The SET TOCLOADRETENTION command can be used to specify the approximate number of minutes that an unreferenced table of contents (TOC) remains loaded in the IBM Spectrum Protect™ database. The IBM Spectrum Protect server-wide TOC retention value determines how long a loaded TOC is retained in the database after the latest access to information in the TOC.

Because TOC information is loaded into temporary database tables, this information is lost if the server is halted, even if the TOC retention period did not elapse. At installation, the retention time is set to 120 minutes. Use the QUERY STATUS command to see the TOC retention time.

Issue the QUERY NASBACKUP command to display information about the file system image objects that are backed up for a specific NAS node and file space. By issuing the command, you can see a display of all backup images that are generated by NDMP and whether each image has a corresponding TOC.

Tip: The IBM Spectrum Protect server can store a full backup in excess of the number of versions you specified if that full backup has dependent differential backups. Full NAS backups with dependent differential backups behave like other base files with dependent subfiles. Due to the retention time specified in the RETEXTRA setting, the full NAS backup is not expired, and the version is displayed in the output of a QUERY NASBACKUP command. For information about setting data retention policies, see Customizing policies.

Use the QUERY TOC command to display files and directories in a backup image that is generated by NDMP. By issuing the QUERY TOC server command, you can display all directories and files within a single specified TOC. The specified TOC is accessed in a storage pool each time the QUERY TOC command is issued because this command does not load TOC information into the IBM Spectrum Protect database. Then, use the RESTORE NODE command with the FILELIST parameter to restore individual files.

Preventing long-running, inactive NDMP connections from closing

To prevent firewalls from closing NDMP connections that are long-running but inactive, you can enable Transmission Control Protocol (TCP) keepalive on the NDMP control connections.

About this task

The IBM Spectrum Protect™ server initiates control connections to NAS devices during NDMP backup or restore operations. These control connections might remain open and inactive for an extended amount of time. For example, suppose that two NDMP operations are started for the same NAS device. The control connection for one NDMP operation might remain open but inactive if the operation requires a resource, for example, a tape drive or sequential volume, that is being used by the other NDMP operation.

Some firewall software is configured to automatically close network connections that are inactive for a specified length of time. If a firewall exists between an IBM Spectrum Protect server and a NAS device, it is possible that the firewall can close NDMP control connections unexpectedly and cause the NDMP operation to fail.

The IBM Spectrum Protect server provides a mechanism, TCP keepalive, that you can enable to prevent long-running, inactive connections from being closed. If TCP keepalive is enabled, small packets are sent across the network at predefined intervals to the connection partner.

Restriction: To prevent errors, do not enable TCP keepalive in certain types of environments. One example is environments that do not have firewalls between the IBM Spectrum Protect server and a NAS device. Another example is environments with firewalls that tolerate long-running, inactive connections. Enabling TCP keepalive in these types of environments can cause an idle connection to be inadvertently closed if the connection partner temporarily fails to respond to TCP keepalive packets.

- Enabling TCP keepalive
To enable TCP keepalive, which keeps NDMP connections open, use the NDMPENABLEKEEPALIVE server option.

- **AIX** | **Linux** | **Windows** Specifying connection idle time for TCP keepalive
To specify the amount of connection idle time, in minutes, before the first TCP keepalive packet is sent, use the NDMPKEEPIDLEMINUTES server option.

Enabling TCP keepalive

To enable TCP keepalive, which keeps NDMP connections open, use the NDMPENABLEKEEPALIVE server option.

Procedure

Add the option to the server options file dsmserv.opt:

```
ndmpenablekeepalive yes
```

Related reference:

NDMPENABLEKEEPALIVE

AIX | **Linux** | **Windows**

Specifying connection idle time for TCP keepalive

To specify the amount of connection idle time, in minutes, before the first TCP keepalive packet is sent, use the NDMPKEEPIDLEMINUTES server option.

Procedure

Add the option to the server options file dsmserv.opt:

```
ndmpkeepidleminutes minutes
```

Related reference:

NDMPKEEPIDLEMINUTES

Configuring IBM Spectrum Protect for NDMP operations

You can configure IBM Spectrum Protect™ to back up and recover data on NAS file servers by using NDMP. The configuration procedure differs depending on whether you plan to back up data from a nonclustered or clustered NAS file server.

Before you begin

Review the restrictions on NDMP backup operations:

- Data deduplication is supported only with NetApp file servers that are not using the SnapMirror to Tape feature.
- Container storage pools are supported only with NetApp file servers that are not using the SnapMirror to Tape feature.

For more information about the types of storage pools that are supported for different brands of file servers, see Storage pool types for NDMP operations.

- Configuring IBM Spectrum Protect for NDMP operations in a nonclustered environment
To use IBM Spectrum Protect for network data management protocol (NDMP) operations in a nonclustered environment, you must set up the tape library and media and complete additional configuration steps.
- Configuring IBM Spectrum Protect for NDMP operations in a NetApp clustered environment
You can back up data from a NetApp cluster to a directly attached tape device or to an IBM Spectrum Protect server, which stores the data in a storage pool. You can back up the entire cluster to a single IBM Spectrum Protect node or parts of the cluster to multiple nodes.

Configuring IBM Spectrum Protect for NDMP operations in a nonclustered environment

To use IBM Spectrum Protect™ for network data management protocol (NDMP) operations in a nonclustered environment, you must set up the tape library and media and complete additional configuration steps.

Before you begin

1. Review the restrictions on NDMP operations:
 - o To configure IBM Spectrum Protect for NDMP operations in a nonclustered environment, you must use a network-attached storage (NAS) device that is validated for use with IBM Spectrum Protect through the Ready for IBM® validation program.
 - o Data deduplication operations and container storage pools are supported only with NetApp file servers that do not use the SnapMirror to Tape feature. All other NAS devices that are validated for use with IBM Spectrum Protect in accordance with the Ready for IBM validation program must use non-deduplicated, non-container storage pools. For more information about the types of storage pools that are supported for different brands of file servers, see Storage pool types for NDMP operations.
2. Register the IBM Spectrum Protect Extended Edition license. To back up and restore NAS file server data by using the IBM Spectrum Protect server, IBM Spectrum Protect Extended Edition is required.

Procedure

1. Set up the tape library and media. See Configuring a tape library for NDMP operations, where the following steps are described in more detail.
 - a. Attach the SCSI or virtual tape library (VTL) library to the NAS file server or to the IBM Spectrum Protect server, or attach the ACSLS library or 349X library to the IBM Spectrum Protect server.
 - b. Define the library with a library type of SCSI, VTL, ACSLS, or 349X.
 - c. Define a device class for the tape drives.
 - d. Define a storage pool for NAS backup media.
 - e. Optional: Define a storage pool for storing a table of contents.
2. Configure IBM Spectrum Protect policy for managing NAS image backups. See Configuring an IBM Spectrum Protect policy for NDMP operations.
3. Register a NAS file server node with the IBM Spectrum Protect server. See Registering NAS nodes with the IBM Spectrum Protect server.
4. Define a data mover for the NAS file server. See Defining a data mover for a NAS file server.
5. Define a path from either the IBM Spectrum Protect server or the NAS file server to the library. See Defining paths to libraries for NDMP operations.
6. Define the tape drives to IBM Spectrum Protect, and define the paths to those drives from the NAS file server and optionally from the IBM Spectrum Protect server. See Defining paths for NDMP operations.
7. Check tapes into the library and label them.

AIX | **Linux** Tape volumes must be labeled before the server can use them. You can use the LABEL LIBVOLUME command, or you can use the AUTOLABEL parameter with the DEFINE LIBRARY and UPDATE LIBRARY commands.

Windows All media must be labeled. Labeling media with an automated library requires you to check media into the library. To label volumes with the LABEL LIBVOLUME command, specify the CHECKIN parameter. To automatically label tape volumes in SCSI-type libraries, use the AUTOLABEL parameter on the DEFINE LIBRARY and UPDATE LIBRARY commands.

For instructions, see LABEL LIBVOLUME, DEFINE LIBRARY, and UPDATE LIBRARY.

8. Optional: Set up scheduled backups for NAS file servers. See Scheduling NDMP operations.
 9. Optional: Define a virtual file space name. See Defining virtual file spaces.
 10. Optional: Configure the tape-to-tape copy function to back up data. See Backing up data with the tape-to-tape function.
 11. Optional: Configure the tape-to-tape copy function to move data to a different tape technology. See Moving data with the tape-to-tape copy function.
- Configuring an IBM Spectrum Protect policy for NDMP operations
With policies, you can manage the number and retention time of NDMP image backup versions.
 - Tape libraries and drives for NDMP operations
Most of the planning that is required to implement backup and recovery operations that use NDMP is related to device configuration. You have choices about how to connect and use the libraries and drives.
 - Attaching tape library robotics for NAS-attached libraries
If you plan to back up NAS data to a library that is directly attached to the NAS device and use a SCSI tape library, you must determine where to attach the library.
 - Registering NAS nodes with the IBM Spectrum Protect server
Register the NAS file server as an IBM Spectrum Protect node, specifying TYPE=NAS. This node name is used to track the image backups for the NAS file server.
 - Defining a data mover for a NAS file server
Define a data mover for each NAS file server by using NDMP operations in your environment. The data mover name must

match the node name that you specified when you registered the NAS node to the IBM Spectrum Protect server.

- Defining paths for NDMP operations
For NDMP operations, you create paths to drives and to libraries.
- Scheduling NDMP operations
You can schedule backup or restore operations for images that are produced by NDMP operations. Use administrative schedules that process the BACKUP NODE or RESTORE NODE administrative commands.
- Defining virtual file spaces
Use a virtual file space definition to complete NAS directory-level backups. To reduce backup and restore times for large file systems, map a directory path from a NAS file server to a virtual file space name on the IBM Spectrum Protect server.
- Backing up data with the tape-to-tape function
When you use the NDMP tape-to-tape function to back up data, the library type can be SCSI, 349X, or ACSLS (automated cartridge system library software). Drives can be shared between the NAS devices and the IBM Spectrum Protect server.
- Moving data with the tape-to-tape copy function
To move data from a previous tape technology to a new tape technology by using the NDMP tape-to-tape copy operation, you must complete the standard steps in your configuration setup and additional steps.

Configuring an IBM Spectrum Protect policy for NDMP operations

With policies, you can manage the number and retention time of NDMP image backup versions.

About this task

For more information, see Policies for backups initiated with an IBM Spectrum Protect server.

Procedure

Complete the following steps to configure a policy for NDMP operations:

1. Create a policy domain for NAS (network-attached storage) file servers. For example, to define a policy domain that is named NASDOMAIN, enter the following command:

```
define domain nasdomain description='Policy domain for NAS file servers'
```
2. Create a policy set in that domain. For example, to define a policy set named STANDARD in the policy domain that is named NASDOMAIN, issue the following command:

```
define policysset nasdomain standard
```
3. Define a management class, and then assign the management class as the default for the policy set. For example, to define a management class that is named MC1 in the STANDARD policy set, and assign it as the default, issue the following commands:

```
define mgmtclass nasdomain standard mcl  
assign defmgmtclass nasdomain standard mcl
```
4. Define a backup copy group in the default management class. The destination must be the storage pool that you created for backup images that are produced by NDMP operations. In addition, you can specify the number of backup versions to retain. For example, to define a backup copy group for the MC1 management class where up to four versions of each file system are retained in the storage pool that is named NASPOOL, issue the following command:

```
define copygroup nasdomain standard mcl destination=naspool verexists=4
```

If you want to create a table of contents for your backups, the TOCDESTINATION parameter of the copy group must contain the name of the primary storage pool.

```
define copygroup nasdomain standard mcl destination=naspool  
tocdestination=tocpool verexists=4
```

Important: When you define a copy group for a management class to which a file system image produced by NDMP is bound, be sure that the DESTINATION parameter specifies the name of a storage pool that is defined for NDMP operations. If the DESTINATION parameter specifies an invalid storage pool, backups by NDMP fail.

5. Activate the policy set. For example, to activate the STANDARD policy set in the NASDOMAIN policy domain, issue the following command:

```
activate policyset nasdomain standard
```

The policy is ready to be used. Nodes are associated with a policy when they are registered. For more information, see [Registering NAS nodes with the IBM Spectrum Protect server](#).

- Policies for backups initiated with an IBM Spectrum Protect server
You can register a network-attached storage (NAS) file server as a node by using network data management protocol (NDMP) operations. Under the direction of the IBM Spectrum Protect server, the NAS file server backs up and restores a file system and directory images to a tape library.
- Policies for backups initiated with the client interface
When a client node initiates a backup, the policy is affected by the option file for that client node.
- Determination of the NAS backup location
When IBM Spectrum Protect uses NDMP to protect NAS file servers, the IBM Spectrum Protect server controls operations. During this time, the NAS file server transfers the data, either to an attached library or directly to the IBM Spectrum Protect server.

Policies for backups initiated with an IBM Spectrum Protect server

You can register a network-attached storage (NAS) file server as a node by using network data management protocol (NDMP) operations. Under the direction of the IBM Spectrum Protect™ server, the NAS file server backs up and restores a file system and directory images to a tape library.

The IBM Spectrum Protect server initiates the backup, allocates a drive, and selects and mounts the media. The NAS file server then transfers the data to tape.

Because the NAS file server backs up the data, the data is stored in its own format. For most NAS file servers, the data is stored in the NDMPDUMP data format. For NetApp file servers, the data is stored in the NETAPPDUMP data format. For EMC file servers, the data is stored in the CELERRADUMP data format. To manage NAS file server image backups, copy groups for NAS nodes must point to a storage pool that has a data format of NDMPDUMP, NETAPPDUMP, or CELERRADUMP.

The following backup copy group attributes are ignored for NAS images:

- Frequency
- Mode
- Retain Only Versions
- Serialization
- Versions Data Deleted

To set up the required policy for NAS nodes, you can define a new, separate policy domain.

When the IBM Spectrum Protect server creates a table of contents (TOC), you can view a collection of individual files and directories that are backed up by using NDMP. Then, you can select which file and directories to restore. To establish where to send data and store the table of contents, set the policy in the following way:

- Ensure that image backup data is sent to a storage pool with an NDMPDUMP, NETAPPDUMP, or CELERRADUMP format.
- Ensure that the table of contents is sent to a storage pool with a NATIVE or NONBLOCK format.

Policies for backups initiated with the client interface

When a client node initiates a backup, the policy is affected by the option file for that client node.

You can control the management classes that are applied to backup images produced by NDMP (network data management protocol) operations regardless of which node initiates the backup. You can complete this task by creating a set of options to be used by the client nodes. The option set can include an `include.fs.nas` statement to specify the management class for NAS (network-attached storage) file server backups.

Tip: You can define an option set by using the `DEFINE CLOPTSET` command. Then, add a client option to the option set by using the `DEFINE CLIENTOPT` command. You can assign an option set to a client by completing the following steps:

1. Open the Operations Center Overview page and click Clients.
2. Double-click the client and click Properties.
3. In the Option set field, select an option set and click Save.

For instructions about using the DEFINE CLOPTSET command, see DEFINE CLOPTSET (Define a client option set name). For instructions about using the DEFINE CLIENTOPT command, see DEFINE CLIENTOPT (Define an option to an option set).

Determination of the NAS backup location

When IBM Spectrum Protect™ uses NDMP to protect NAS file servers, the IBM Spectrum Protect server controls operations. During this time, the NAS file server transfers the data, either to an attached library or directly to the IBM Spectrum Protect server.

You can also use a backup-archive client to back up a NAS file server by mounting the NAS file system on the client computer and then backing up as usual. You can use either a Network File System (NFS) mount or a Common Internet File System (CIFS) map.

For a description of the backup-and-restore methods, see Table 1.

Tip: You can use a single method or a combination of methods in your individual storage environment.

Table 1. Comparing methods for backing up NDMP data

Property	NDMP: File server to server	NDMP: File server to attached library	Backup-archive client to server
Network data traffic	All backup data goes across the LAN from the NAS file server to the server.	The server controls operations remotely, but the NAS device moves the data locally.	All backup data goes across the LAN from the NAS device to the client and then to the server.
File server processing during backup	Less file server processing is required, compared to the backup-archive client method, because the backup does not use file access protocols such as NFS and CIFS.	Less file server processing is required, compared to the backup-archive client method, because the backup does not use file access protocols such as NFS and CIFS.	File backup operations require more server processing resources for file access protocols such as NFS and CIFS.
Distance between devices	The IBM Spectrum Protect server must be within SCSI or Fibre Channel range of the tape library.	The IBM Spectrum Protect server can be distant from the NAS file server and the tape library.	The IBM Spectrum Protect server must be within SCSI or Fibre Channel range of the tape library.
Firewall considerations	More stringent than filer-to-attached-library because communications can be initiated by either the IBM Spectrum Protect server or the NAS file server.	Less stringent than filer-to-server because communications can be initiated only by the IBM Spectrum Protect server.	Client passwords and data are encrypted.
Security considerations	Data is sent decrypted from a NAS file server to an IBM Spectrum Protect server.	This method must be used in a trusted environment because port numbers are not secure.	Port number configuration allows for secure administrative sessions within a private network.
Load on the IBM Spectrum Protect server	Higher CPU workload is required to manage all back-end data processes (for example, migration).	CPU workload is reduced because migration and reclamation are not supported.	Higher CPU workload is required to manage all back-end data processes.
Backup of primary storage pools to copy storage pools	Data can be backed up only to copy storage pools that have the NATIVE data format.	Data can be backed up only to copy storage pools that have the same NDMP data format (NETAPPDUMP, CELERRADUMP, or NDMPDUMP).	Data can be backed up only to copy storage pools that have the NATIVE data format.
Restore of primary storage pools and volumes from copy storage pools	Data can be restored only to storage pools and volumes that have the NATIVE data format.	Data can be restored only to storage pools and volumes that have the same NDMP data format.	Data can be restored only to storage pools and volumes that have the NATIVE data format.

Property	NDMP: File server to server	NDMP: File server to attached library	Backup-archive client to server
Moving NDMP data from storage pool volumes	Data can be moved to another storage pool only if it has a NATIVE data format.	Data can be moved to another storage pool only if it has the same NDMP data format.	Data can be moved to another storage pool only if it has a NATIVE data format.
Migration from one primary storage pool to another	Supported	Not supported	Supported
Reclamation of a storage pool	Supported	Not supported	Supported
Simultaneous-write operations during backups	Not supported	Not supported	Supported
Export and import operations	Not supported	Not supported	Supported
Backup set generation	Not supported	Not supported	Supported
Cyclic Redundancy Checking (CRC) when data is moved by using IBM Spectrum Protect processes	Supported	Not supported	Supported
Validation by using IBM Spectrum Protect audit commands	Supported	Not supported	Supported
Disaster recovery manager	Supported	Supported	Supported

Tape libraries and drives for NDMP operations

Most of the planning that is required to implement backup and recovery operations that use NDMP is related to device configuration. You have choices about how to connect and use the libraries and drives.

Many of the configuration choices you have for libraries and drives are determined by the hardware features of your libraries. You can set up NDMP operations with any supported library and drives. However, the more features your library has, the more flexibility you can exercise in your implementation.

You might start by answering the following questions:

- What type of library (SCSI, ACSLS, or 349X) will you use?
 - If you are using a SCSI library, do you want to attach tape library robotics to the IBM Spectrum Protect™ server or to the network-attached storage (NAS) file server?
 - Will you want to move your NDMP data to tape?
 - How do you want to use the tape drives in the library?
 - Dedicate all tape drives to NDMP operations.
 - Dedicate some tape drives to NDMP operations and others to traditional IBM Spectrum Protect operations.
 - Share tape drives between NDMP operations and traditional IBM Spectrum Protect operations.
 - Will you back up data tape-to-tape for disaster recovery functions?
 - Will you send backup data to a single IBM Spectrum Protect server instead of attaching a tape library to each NAS device?
 - Do you want to keep all hardware on the IBM Spectrum Protect server and send NDMP data over the LAN?
-
- Determining library drive usage when backing up to NAS-attached libraries
Drives can be used for multiple purposes because of the flexible configurations that are allowed by IBM Spectrum Protect. For NDMP operations, the NAS file server must have access to the drive. The IBM Spectrum Protect server can also have access to the same drive, depending on your hardware connections and limitations.
 - Configuring a tape library for NDMP operations
You can configure a tape library to back up a network-attached storage (NAS) device to tape.


Determining library drive usage when backing up to NAS-attached libraries

Drives can be used for multiple purposes because of the flexible configurations that are allowed by IBM Spectrum Protect™. For NDMP operations, the NAS file server must have access to the drive. The IBM Spectrum Protect server can also have access to the same drive, depending on your hardware connections and limitations.

About this task

All drives are defined to the IBM Spectrum Protect server. However, the same drive can be defined for both traditional IBM Spectrum Protect operations and NDMP operations. Figure 1 illustrates one possible configuration. The IBM Spectrum Protect server has access to drives 2 and 3, and each NAS file server has access to drives 1 and 2.

Figure 1. IBM Spectrum Protect drive usage example



To create the configuration that is shown in Figure 1, complete the following steps:

Procedure

1. Define all three drives to IBM Spectrum Protect.
2. Define paths from the IBM Spectrum Protect server to drives 2 and 3. Because drive 1 is not accessed by the server, no path is defined.
3. Define each NAS file server as a separate data mover.
4. Define paths from each data mover to drive 1 and to drive 2.

Results

To use the IBM Spectrum Protect back-end data movement operations, the IBM Spectrum Protect server requires two available drive paths from a single NAS data mover. The drives can be in different libraries and can have different device types that are supported by NDMP. You can make copies between two different tape devices. For example, the source tape drive can be a DLT drive in a library and the target drive can be an LTO drive in another library.

During IBM Spectrum Protect back-end data movements, the IBM Spectrum Protect server locates a NAS data mover that supports the same data format as the data to be copied from and that has two available mount points and paths to the drives. If the IBM Spectrum Protect server cannot locate such a data mover, the requested data movement operation is not performed. The number of available mount points and drives depends on the mount limits of the device classes for the storage pools that are involved in the back-end data movements.

If the back-end data movement function supports multiprocessing, each concurrent IBM Spectrum Protect back-end data movement process requires two available mount points and two available drives. To run two IBM Spectrum Protect processes concurrently, at least four mount points and four drives must be available.

For more information, see [Defining paths for NDMP operations](#).

Configuring a tape library for NDMP operations

You can configure a tape library to back up a network-attached storage (NAS) device to tape.

Procedure

Complete the following steps to set up tape libraries for NDMP operations:

1. Connect the library and drives to be used for NDMP operations.
 - a. Connect the SCSI library. Before you set up a SCSI tape library for NDMP operations, determine whether you want to attach your library robotics control to the IBM Spectrum Protect™ server or to the NAS file server. See Tape libraries and drives for NDMP operations. Connect the SCSI tape library robotics to the IBM Spectrum Protect server or to the NAS file server. Refer to your device manufacturer documentation for instructions.

If the library is connected to IBM Spectrum Protect, make a SCSI or Fibre Channel connection between the IBM Spectrum Protect server and the library robotics control port. Then, connect the NAS file server with the drives.

If the library is connected to the NAS file server, make a SCSI or Fibre Channel connection between the NAS file server and the library robotics and drives.

- b. Connect the ACSLS Library. Connect the ACSLS tape library to the IBM Spectrum Protect server.
 - c. Connect the 349X Library. Connect the 349X tape library to the IBM Spectrum Protect server.
2. Define the library for your library device by issuing the DEFINE LIBRARY command. The library must be a single device type, not a mixed device. Issue one of the following commands to define the library depending on the type of device that you are configuring:

SCSI Library

```
define library tsmlib libtype=scsi
```

ACSLs Library

```
define library acslib libtype=acsls acsid=1
```

349X Library

```
define library tsmlib libtype=349x
```

3. Define a device class for your NDMP device by issuing the DEFINE DEVCLASS command.

Tip: A device class that is defined with a device type of NAS is not explicitly associated with a specific drive type, for example, LTO. However, the preferred practice is to define a separate device class for different drive types.

In the DEFINE DEVCLASS command, use the following parameters and values:

 - o Specify `DEVTYPE=NAS`.
 - o Specify `MOUNTRETENTION=0`. It is required for NDMP operations.
 - o Specify a value for the `ESTCAPACITY` parameter.

For example, to define a device class that is named NASCLASS for a library that is named NASLIB with an estimated capacity is 40 GB for the media, issue the following command:

```
define devclass nasclass devtype=nas library=naslib mountretention=0  
estcapacity=40g
```

4. Define a storage pool for NDMP media by issuing the DEFINE STGPOOL command. When NETAPPDUMP, CELERRADUMP, or NDMPDUMP is designated as the type of storage pool, managing the storage pools that are produced by NDMP operations is different from managing storage pools that contain media for traditional IBM Spectrum Protect backups. IBM Spectrum Protect operations use storage pools that are defined with a NATIVE or NONBLOCK data format. If you select NETAPPDUMP, CELERRADUMP, or NDMPDUMP, NDMP operations require storage pools with a data format that matches the NAS file server and the selected backup method. Maintaining separate storage pools for data from different NAS vendors is optimal, even though the data format for both is NDMPDUMP. For example, to define a storage pool that is named NDMPPPOOL for a file server, which is not a NetApp or a Celerra file server, issue the following command:

```
define stgpool ndmppool nasclass maxscratch=10 dataformat=ndmpdump
```

To define a storage pool that is named NASPOOL for a NetApp file server, issue the following command:

```
define stgpool naspool nasclass maxscratch=10 dataformat=netappdump
```

To define a storage pool that is named CELERRAPOOL for an EMC Celerra file server, issue the following command:

```
define stgpool celerrapool nasclass maxscratch=10 dataformat=celerradump
```

Attention: Ensure that you do not accidentally use storage pools that are defined for NDMP operations in traditional IBM Spectrum Protect operations. Be especially careful when you assign the storage pool name as the value for the DESTINATION parameter of the DEFINE COPYGROUP command. Unless the destination is a storage pool with the appropriate data format, the backup can fail.

- Optional: Define a storage pool for a table of contents. If you plan to create a table of contents, you must also define a disk storage pool in which to store the table of contents. You must set up policy so that the IBM Spectrum Protect server stores the table of contents in a different storage pool from the one where the backup image is stored. The table of contents is treated like any other object in that storage pool. For example, to define a storage pool that is named TOCPool for a DISK device class, issue the following command:

```
define stgpool tocpool disk
```

Then define volumes for the storage pool.

AIX | **Linux** For more information about defining volumes, see [Configuring random access volumes on disk devices \(V7.1.1\)](#).

Windows For more information about defining volumes, see [Configuring random access volumes on disk devices \(V7.1.1\)](#).

AIX | **Linux** For more information about configuring libraries, see [Configuring libraries for use by a server](#).

Related reference:

DEFINE DEVCLASS (Define a device class)

Attaching tape library robotics for NAS-attached libraries

If you plan to back up NAS data to a library that is directly attached to the NAS device and use a SCSI tape library, you must determine where to attach the library.

About this task

You must determine whether to attach the library robotics to the IBM Spectrum Protect™ server or to the NAS file server. Regardless of where you connect library robotics, tape drives must always be connected to the NAS file server for NDMP operations.

Distance and your available hardware connections are factors to consider for SCSI libraries. If the library does not have separate ports for robotics control and drive access, the library must be attached to the NAS file server because the NAS file server must have access to the drives. If your SCSI library has separate ports for robotics control and drive access, you can choose to attach the library robotics to either the IBM Spectrum Protect server or the NAS file server. If the NAS file server is at a different location from the IBM Spectrum Protect server, the distance might mean that you must attach the library to the NAS file server.

Whether you are using a SCSI, ACSLS, or 349X library, you have the option of dedicating the library to NDMP operations, or of using the library for NDMP operations. You can also use the library for most traditional IBM Spectrum Protect operations.

Table 1. Summary of configurations for NDMP operations

Configuration	Distance between IBM Spectrum Protect server and library	Library sharing	Drive sharing between IBM Spectrum Protect and NAS file server	Drive sharing between NAS file servers	Drive sharing between storage agent and NAS file server
Configuration 1 (SCSI library that is connected to the IBM Spectrum Protect server)	Limited by SCSI or FC connection	Supported	Supported	Supported	Supported
Configuration 2 (SCSI library that is connected to the NAS file server)	No limitation	Not supported	Supported	Supported	Not supported
Configuration 3 (349X library)	Might be limited by 349X connection	Supported	Supported	Supported	Supported
AIX Windows Configuration 4 (ACSLs library)	AIX Windows Might be limited by ACSLS connection	AIX Windows Supported	AIX Windows Supported	AIX Windows Supported	AIX Windows Supported

- Configuration 1: SCSI library connected to the IBM Spectrum Protect server
In this configuration, the tape library must have separate ports for robotics control and for drive access. In addition, the library must be within Fibre Channel range or SCSI bus range of both the IBM Spectrum Protect server and the NAS file server.
- Configuration 2: SCSI library connected to the NAS file server
In this configuration, the library robotics and the drives must be physically connected directly to the NAS file server. Paths must be defined from the NAS data mover to the library and drives. No physical connection is required between the IBM Spectrum Protect server and the SCSI library.
- Configuration 3: 349x library connected to the IBM Spectrum Protect server
For this configuration, you connect the tape library to the system as for traditional operations.
- Configuration 4: ACSLS library connected to the IBM Spectrum Protect server
For this configuration, connect the tape library to the system as you do for traditional IBM Spectrum Protect operations.


Configuration 1: SCSI library connected to the IBM Spectrum Protect server

In this configuration, the tape library must have separate ports for robotics control and for drive access. In addition, the library must be within Fibre Channel range or SCSI bus range of both the IBM Spectrum Protect™ server and the NAS file server.

In this configuration, the IBM Spectrum Protect server controls the SCSI library through a direct, physical connection to the library robotics control port. For NDMP operations, the drives in the library are connected directly to the NAS file server, and a path must be defined from the NAS data mover to each of the drives to be used. The NAS file server transfers data to the tape drive at the request of the IBM Spectrum Protect server. To also use the drives for IBM Spectrum Protect operations, connect the IBM Spectrum Protect server to the tape drives and define paths from the server to the tape drives.

This configuration also supports an IBM Spectrum Protect storage agent having access to the drives for its LAN-free operations, and the IBM Spectrum Protect server can be a library manager.

Figure 1. Configuration 1: SCSI library connected to IBM Spectrum Protect server

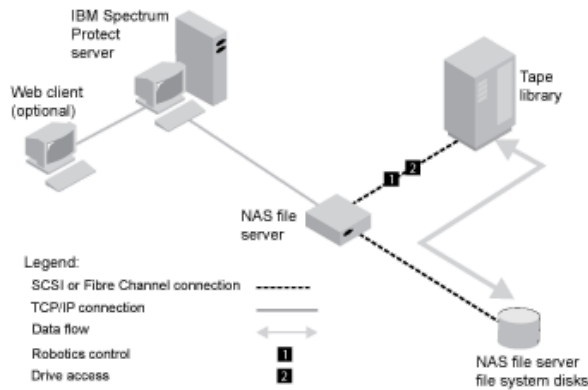


Configuration 2: SCSI library connected to the NAS file server

In this configuration, the library robotics and the drives must be physically connected directly to the NAS file server. Paths must be defined from the NAS data mover to the library and drives. No physical connection is required between the IBM Spectrum Protect™ server and the SCSI library.

The IBM Spectrum Protect server controls library robotics by sending library commands across the network to the NAS file server. The NAS file server passes the commands to the tape library. Any responses that are generated by the library are sent to the NAS file server, and passed back across the network to the IBM Spectrum Protect server. This configuration supports a physically distant IBM Spectrum Protect server and NAS file server. For example, the IBM Spectrum Protect server is in one city, while the NAS file server and tape library are in another city.

Figure 1. Configuration 2: SCSI library connected to the NAS file server



Configuration 3: 349x library connected to the IBM Spectrum Protect server


For this configuration, you connect the tape library to the system as for traditional operations.

In this configuration, the 349X tape library is controlled by the IBM Spectrum Protect™ server. The IBM Spectrum Protect server controls the library by passing the request to the 349X library manager through TCP/IP.

To complete NAS (network-attached storage) backup or restore operations, the NAS file server must be able to access one or more tape drives in the 349X library. Any tape drives used for NAS operations must be physically connected to the NAS file server, and paths need to be defined from the NAS data mover to the drives. The NAS file server transfers data to the tape drive at the request of the IBM Spectrum Protect server. Follow the manufacturer's instructions to attach the device to the server system.

This configuration supports a physically distant IBM Spectrum Protect server and NAS file server. For example, the IBM Spectrum Protect server might be in one city, while the NAS file server and tape library are in another city.

Figure 1. Configuration 3: 349x library connected to the IBM Spectrum Protect server



Related information:

[Attaching devices for the server](#)

Configuration 4: ACSLS library connected to the IBM Spectrum Protect server

For this configuration, connect the tape library to the system as you do for traditional IBM Spectrum Protect™ operations.

The ACSLS (automated cartridge system library software) tape library is controlled by the IBM Spectrum Protect server. The IBM Spectrum Protect server controls the library by passing the request to the ACSLS library server through TCP/IP. The ACSLS library supports library sharing and LAN-free operations.


Windows Restriction: To use ACSLS functions, StorageTek Library Attach software must be installed. For more information, see ACSLS-managed libraries (V7.1.1).

To complete NAS (network-attached storage) backup or restore operations, the NAS file server must be able to access one or more tape drives in the ACSLS library. Any tape drives used for NAS operations must be physically connected to the NAS file server, and any paths need to be defined from the NAS data mover to the drives. The NAS file server transfers data to the tape drive at the request of the IBM Spectrum Protect server. Follow the manufacturer's instructions to attach the device to the server system.

This configuration supports a physically distant IBM Spectrum Protect server and NAS file server. For example, the IBM Spectrum Protect server might be in one city while the NAS file server and tape library are in another city.

To also use the drives for IBM Spectrum Protect operations, connect the IBM Spectrum Protect server to the tape drives and define paths from the IBM Spectrum Protect server to the tape drives.

Figure 1. Configuration 4: ACSLS library connected to the IBM Spectrum Protect server



Related information:

[Attaching devices for the server](#)

Registering NAS nodes with the IBM Spectrum Protect server

Register the NAS file server as an IBM Spectrum Protect™ node, specifying TYPE=NAS. This node name is used to track the image backups for the NAS file server.

Procedure

To register a NAS file server as a node named NASNODE1, with a password of NASPWD1, in a policy domain that is named NASDOMAIN, issue the following example command:

```
register node nasnode1 naspwd1 domain=nasdomain type=nas
```

If you are using a client option set, specify the option set when you register the node. You can verify that this node is registered by issuing the following command:

```
query node type=nas
```

Remember: You must specify TYPE=NAS so that only NAS nodes are displayed.

Defining a data mover for a NAS file server

Define a data mover for each NAS file server by using NDMP operations in your environment. The data mover name must match the node name that you specified when you registered the NAS node to the IBM Spectrum Protect™ server.

About this task

The definition for a NAS data mover contains the network address, authorization, and data formats that are required for NDMP operations. A data mover enables communication and ensures authority for NDMP operations between the IBM Spectrum Protect server and the NAS file server.

Procedure

To define a data mover, use the DEFINE DATAMOVER command.

Example

For example, define a data mover with these parameters:

- The NAS node is named NASNODE1.
- The high-level address is an IP address for the NAS file server, either a numerical address or a host name.
- The low-level address is the IP port for NDMP sessions with the NAS file server. The default is port number 10000.
- The user ID is the ID defined to the NAS file server that authorizes an NDMP session with the NAS file server. For this example, the user ID is the administrative ID for the NetApp file server.
- The password parameter is a valid password for authentication to an NDMP session with the NAS file server.
- The data format is NETAPPDUMP. This is the data format that the NetApp file server uses for tape backup. This data format must match the data format of the target storage pool.

Enter the following command:

```
define datamover nasnode1 type=nas haddress=netapp2 lladdress=10000 userid=root  
password=admin dataformat=netappdump
```

Related reference:

DEFINE DATAMOVER (Define a data mover)

Defining paths for NDMP operations

For NDMP operations, you create paths to drives and to libraries.

- Defining paths to drives for NDMP operations
The method that you choose for creating paths to drives depends on whether the drives are accessed by a NAS file server and the IBM Spectrum Protect server or only by a NAS file server.
- Defining paths to libraries for NDMP operations
Define a path to the SCSI library from either the IBM Spectrum Protect server or the NAS file server.

Defining paths to drives for NDMP operations

The method that you choose for creating paths to drives depends on whether the drives are accessed by a NAS file server and the IBM Spectrum Protect™ server or only by a NAS file server.

- Defining paths for drives attached to a NAS file server and to the IBM Spectrum Protect server
If a tape drive must be accessed by a network-attached storage (NAS) file server and the IBM Spectrum Protect server, you must create two paths. One path exists between the tape drive and the NAS file server. The other path exists between the tape drive and the IBM Spectrum Protect server.
- Defining paths for drives attached only to NAS file servers
If a tape drive must be accessed only by a NAS file server and not by the IBM Spectrum Protect server, only a single path between the tape drive and the NAS file server is required.
- Obtaining names for devices attached to NAS file servers
For paths from a NAS data mover, the value of the DEVICE parameter in the DEFINE PATH command is the name by which the NAS file server knows a library or drive.

Defining paths for drives attached to a NAS file server and to the IBM Spectrum Protect server

If a tape drive must be accessed by a network-attached storage (NAS) file server and the IBM Spectrum Protect™ server, you must create two paths. One path exists between the tape drive and the NAS file server. The other path exists between the tape drive

and the IBM Spectrum Protect server.

Procedure

Complete the following steps:

1. If the drive is not defined for the IBM Spectrum Protect server, create the drive definition. For example, to define a drive NASDRIVE1 for a library NASLIB, issue the following command:

```
define drive naslib nasdrive1 element=autodetect
```

Remember: If the drive is attached to the IBM Spectrum Protect server, the element address is automatically detected.

2. Map the NAS drive name to the corresponding drive definition on the IBM Spectrum Protect server:
 - o On the IBM Spectrum Protect server, issue the QUERY DRIVE FORMAT=DETAILED command to obtain the worldwide name (WWN) and serial number for the drive that will be connected to the NAS file server.
 - o On the NAS device, obtain the tape device name, serial number, and WWN for the drive.

If the WWN or serial number matches, a drive on a NAS file server is the same as the drive on the IBM Spectrum Protect server.

3. Using the drive name, define a path to the drive from the NAS file server and a path to the drive from the IBM Spectrum Protect server.
 - o For example, to define a path between a tape drive with a device name of rst01 and a NetApp file server, issue the following command:

```
define path nasnode1 nasdrive1 srctype=datamover desttype=drive
  library=naslib device=rst01
```

- o To define a path between the tape drive and the IBM Spectrum Protect server, issue the following command:

AIX

```
define path server1 nasdrive1 srctype=server desttype=drive
  library=naslib device=/dev/rmt0
```

Linux

```
define path server1 nasdrive1 srctype=server desttype=drive
  library=naslib device=/dev/tmscsi/mt0
```

Windows

```
define path server1 nasdrive1 srctype=server desttype=drive
  library=naslib device=mt3.0.0.2
```

Defining paths for drives attached only to NAS file servers

If a tape drive must be accessed only by a NAS file server and not by the IBM Spectrum Protect™ server, only a single path between the tape drive and the NAS file server is required.

Procedure

Complete the following steps:

1. Obtain the SCSI element addresses, worldwide name (WWN), and serial numbers for the drive to be connected to NAS file server.

Restriction: If the SCSI drive is connected only to a NAS file server, the element address is not automatically detected, and you must supply it. If a library has more than one drive, you must specify an element address for each drive.

To obtain a SCSI element address, go to the following device-support websites:

- o **AIX** | **Windows** Supported devices for AIX and Windows
- o **Linux** Supported devices for Linux

Element number assignment and device WWN assignments are also available from tape-library device manufacturers.

2. Create drive definitions by specifying the element addresses identified in the preceding step. Specify the element address in the ELEMENT parameter of the DEFINE DRIVE command. For example, to define a drive NASDRIVE1 with the element address 82 for the library NASLIB, issue the following command:

```
define drive naslib nasdrive1 element=82
```

Attention: For a drive connected only to the NAS file server, do not specify ASNEEDED as the value for the CLEANFREQUENCY parameter of the DEFINE DRIVE command.

3. Obtain the device name, serial number, and WWN for the drive on the NAS device.
4. Using the information that is obtained in steps 1 and 3, map the NAS device name to the element address in the drive definition in the IBM Spectrum Protect server.
5. Define a path between the tape drive and the NAS file server. For example, to define a path between a NetApp file server and a tape drive with a device name of rst01, issue the following command:

```
define path nasnode1 nasdrive1 srctype=datamover desttype=drive
library=naslib device=rst01
```

Obtaining names for devices attached to NAS file servers

For paths from a NAS data mover, the value of the DEVICE parameter in the DEFINE PATH command is the name by which the NAS file server knows a library or drive.

About this task

You can obtain these device names, also known as *special file names*, by querying the NAS file server. For information about how to obtain names for devices that are connected to a NAS file server, see the product information for the file server.

Procedure

- To obtain the device names for tape libraries on a NetApp Release ONTAP 10.0 GX, or later, file server, connect to the file server by using telnet and issue the SYSTEM HARDWARE TAPE LIBRARY SHOW command. To obtain the device names for tape drives on a NetApp Release ONTAP 10.0 GX, or later, file server, connect to the file server by using telnet and issue the SYSTEM HARDWARE TAPE DRIVE SHOW command. For details about these commands, see the NetApp ONTAP GX file server product documentation.
- For releases earlier than NetApp Release ONTAP 10.0 GX, continue to use the SYSCONFIG command. For example, to display the device names for tape libraries, connect to the file server by using telnet and issue the following command:

```
sysconfig -m
```

To display the device names for tape drives, issue the following command:

```
sysconfig -t
```

- For Fibre Channel attached drives and the Celerra data mover, complete the following steps:
 1. Log on to the EMC Celerra control workstation by using an administrative ID. Issue the following command:

```
server_devconfig server_1 -l -s -n
```

Tip: The -l option for this command lists only the device information that was saved in the database of the data mover. The command and option do not display changes to the device configuration that occurred after the last database refresh on the data mover. For details about how to obtain the most recent device configuration for your data mover, see the EMC Celerra documentation.

The output for the server_devconfig command includes the device names for the devices that are attached to the data mover. The device names are listed in the *addr* column, for example:

```
server_1:
Scsi Device Table
name      addr      type  info
tape1     c64t010  tape  IBM ULT3580-TD2 53Y2
ttape1    c96t010  tape  IBM ULT3580-TD2 53Y2
```

2. Map the Celerra device name to the device worldwide name (WWN):
 - a. To list the WWN, log on to the EMC Celerra control workstation and issue the following command. Remember to enter a period (.) as the first character in this command.

```
.server_config server_# -v "fcp bind show"
```

The output for this command includes the WWN, for example:

```
Chain 0064: WWN 500507630f418e29 HBA 2 N_PORT Bound
Chain 0096: WWN 500507630f418e18 HBA 2 N_PORT Bound
```

Tip: The `.server_config` command is an undocumented EMC Celerra command. For more information about how to use it, contact EMC.

- b. Use the chain number to identify the tape device that was listed in the output of the `server_devconfig` command and that has the same WWN, for example:

Tape device name	Chain number	WWN
c64t0l0	0064	500507630f418e29
c96t0l0	0096	500507630f418e18

Celerra commands might behave differently on different EMC Celerra systems and operating system levels. For details, see the EMC Celerra documentation or contact EMC.

Defining paths to libraries for NDMP operations

Define a path to the SCSI library from either the IBM Spectrum Protect™ server or the NAS file server.

Procedure

1. For a SCSI library connected to IBM Spectrum Protect, issue the following example command to define a path from the server, named `SERVER1`, to the SCSI library named `TSMLIB`:

AIX

```
define path server1 tsmlib srctype=server desttype=library
device=/dev/lb1
```

Linux

```
define path server1 tsmlib srctype=server desttype=library
device=/dev/tmscsi/lb1
```

Windows

```
define path server1 tsmlib srctype=server desttype=library
device=lb0.0.0.2
```

2. For a SCSI library connected to a NAS file server, issue the following example command to define a path between a NetApp NAS data mover that is named `NASNODE1` and a library named `NASLIB`:

```
define path nasnode1 naslib srctype=datamover desttype=library device=mc0
```

3. For a 349X library, define a path to the library from the IBM Spectrum Protect server. For example, issue the following command to define a path from the server, named `SERVER1`, to the 349X library named `TSMLIB`:

AIX

```
define path server1 tsmlib srctype=server desttype=library
device=/dev/lmcp0
```

Linux | Windows

```
define path server1 tsmlib srctype=server desttype=library
device=library1
```

Tip: The `DEFINE PATH` command is not required for an automated cartridge system library software (ACSL) library.

Scheduling NDMP operations

You can schedule backup or restore operations for images that are produced by NDMP operations. Use administrative schedules that process the `BACKUP NODE` or `RESTORE NODE` administrative commands.

Procedure

Create an administrative schedule by using the `DEFINE SCHEDULE` command. For example, to create an administrative schedule that is named `NASSCHED` to back up all file systems for node `NASNODE1`, enter the following command:

```
define schedule nassched type=administrative cmd='backup node nasnode1' active=yes
starttime=20:00 period=1 perunits=days
```

The schedule is active, and is set to run at 8 p.m. every day.

Restriction: The BACKUP NODE and RESTORE NODE commands can be used only for nodes of TYPE=NAS.

Related tasks:

🔗 [Tuning the schedule for daily operations](#)

Related reference:

BACKUP NODE (Back up a NAS node)

RESTORE NODE (Restore a NAS node)

DEFINE SCHEDULE (Define a schedule for an administrative command)

Defining virtual file spaces

Use a virtual file space definition to complete NAS directory-level backups. To reduce backup and restore times for large file systems, map a directory path from a NAS file server to a virtual file space name on the IBM Spectrum Protect™ server.

Procedure

To create a virtual file space name for the directory path on the NAS device, issue the DEFINE VIRTUALFSMAPPING command:

```
define virtualfsmapping nas1 /mikesdir /vol/vol1 /mikes
```

This command defines a virtual file space name of /MIKESDIR on the server, which represents the directory path of /VOL/VOL1/MIKES on the NAS file server that is represented by node NAS1. For more information, see Directory-level backup and restore for NDMP operations.

Backing up data with the tape-to-tape function

When you use the NDMP tape-to-tape function to back up data, the library type can be SCSI, 349X, or ACSLS (automated cartridge system library software). Drives can be shared between the NAS devices and the IBM Spectrum Protect™ server.

About this task

When you use the NDMP tape-to-tape copy function, your configuration setup might affect the performance of the IBM Spectrum Protect back-end data movement.

Procedure

To have one NAS device with paths to four drives in a library, use the MOVE DATA command after you complete your configuration setup. This moves data on the volume VOL1 to any available volumes in the same storage pool as VOL1:

```
move data vol1
```

Moving data with the tape-to-tape copy function

To move data from a previous tape technology to a new tape technology by using the NDMP tape-to-tape copy operation, you must complete the standard steps in your configuration setup and additional steps.

About this task

When you use the NDMP tape-to-tape copy function, your configuration setup might affect the performance of the IBM Spectrum Protect™ back-end data movement.

Procedure

In addition to the standard steps in your configuration setup, complete the following steps:

1. Define one drive in the library, lib1, that has previous tape technology:

```
define drive lib1 drv1 element=1035
```

2. Define one drive in the library, lib2, that has new tape technology:

```
define drive lib2 drv1 element=1036
```

3. Define paths from the NAS file server to each drive:

```
define path nas1 drv1 sourcetype=datamover desttype=drive library=lib1 device=rst11  
define path nas1 drv1 sourcetype=datamover desttype=drive library=lib2 device=rst21
```

4. Move data on volume vol1 in the primary storage pool to the volumes in another primary storage pool, nasprimpool2:

```
move data vol1 stgpool=nasprimpool2
```

Configuring IBM Spectrum Protect for NDMP operations in a NetApp clustered environment

You can back up data from a NetApp cluster to a directly attached tape device or to an IBM Spectrum Protect™ server, which stores the data in a storage pool. You can back up the entire cluster to a single IBM Spectrum Protect node or parts of the cluster to multiple nodes.

Before you begin

For an overview of NDMP functionality in IBM Spectrum Protect and NetApp file servers, see technote 7046965. This technote also lists system requirements.

About this task

You can back up data in a NetApp clustered environment to the following storage media:

Tape device that is directly attached to a NAS file server

You can back up data to a tape device that is directly attached to a NAS file server. This is the preferred method. Typically, it is faster to back up data to a directly attached tape device than it is to back up data to an IBM Spectrum Protect storage pool by using a network connection.

Storage pool in the local IBM Spectrum Protect hierarchy

You can back up data to an IBM Spectrum Protect server, which stores the data in a storage pool of type DISK, FILE, or tape. The advantage of storing data in a storage pool is that you can replicate the data for added data protection. You can use existing storage pools or create storage pools. You must have a network connection between the NAS file server and the IBM Spectrum Protect server. The network connection must have sufficient bandwidth to transfer the NAS backup data. Tip: This type of backup is sometimes called a filer-to-server backup.

You can use one of the following backup methods:

Full cluster backup

When you apply this method, the backup data of the entire cluster is owned by a single IBM Spectrum Protect node. Even if you move the volumes within the cluster, full cluster backup operations continue and you do not have to reconfigure backup operations. This is the preferred method.

Partial cluster backup

When you apply this method, you specify a NetApp storage virtual machine (SVM), which determines the scope of the backup operation. The SVM is a virtual server that provides access to part of a cluster. You can specify that each SVM in the cluster backs up data to a separate IBM Spectrum Protect node. This method requires more configuration than the full cluster backup method, and it requires a network connection to transfer data from the SVM to the IBM Spectrum Protect node.

Restriction: You cannot use this method to back up data to a tape device because SVMs do not have direct access to tape devices.

Procedure

1. Select the storage media based on the following questions:

Question	Storage media
Based on your business requirements, is it necessary to back up data to a local tape device?	If the answer is yes, use a directly attached tape device. If the answer is no, use either a directly attached tape device or a local IBM Spectrum Protect storage pool.

Question	Storage media
Does your organization require high-speed backup operations?	If the answer is yes, use a directly attached tape device. If the answer is no, use either a directly attached tape device or a local IBM Spectrum Protect storage pool.
Does your organization have sufficient network bandwidth for NAS backup data?	If the answer is yes, use either a directly attached tape device or a local IBM Spectrum Protect storage pool. If the answer is no, use a directly attached tape device.
Does your organization want to enhance data protection by using replication?	If the answer is yes, use a local IBM Spectrum Protect storage pool. If the answer is no, use either a directly attached tape device or a local IBM Spectrum Protect storage pool.
Are your NAS file servers at remote locations without access to directly attached tape libraries?	If the answer is yes, use a local IBM Spectrum Protect storage pool. If the answer is no, use either a directly attached tape device or a local IBM Spectrum Protect storage pool.

2. Select a backup method based on the following questions:

Question	Backup method
Based on your business requirements, is it necessary to back up data to a directly attached tape device?	If the answer is yes, use the full backup method. If the answer is no, use either the full or partial backup method.
Does your system have sufficient network bandwidth to back up several SVMs without affecting network performance?	If the answer is yes, use either the full or partial backup method. If the answer is no, use the full backup method. The partial backup method might adversely affect system performance.
Are the SVMs distributed across several organizations? For example, are any SVMs controlled by third parties, such as cloud-platform providers?	If the answer is yes, use the partial backup method because SVM owners can control backup operations for individual SVMs. If an SVM owner also owns an IBM Spectrum Protect server, the owner can set up backup operations from the SVM to a server node. In this way, the owner can control the end-to-end process. If the answer is no, use either the full or partial backup method.

3. Configure the system environment based on the storage media and backup method that you chose. Follow the instructions for your selected method:

- Configuring full cluster backups to directly attached tape devices
- Configuring full cluster backups to an IBM Spectrum Protect server
- Configuring partial cluster backups to an IBM Spectrum Protect server

Tip: If you configured IBM Spectrum Protect to back up NetApp clusters by using node-scoped NDMP, consider reconfiguring IBM Spectrum Protect to use NDMP Cluster Aware Backup (CAB). In this way, you can optimize backup operations for NetApp clusters. Follow the instructions in Reconfiguring IBM Spectrum Protect to optimize clustered backups.

- **Configuring full cluster backups to directly attached tape devices**
You can configure IBM Spectrum Protect to back up all volumes in a NetApp cluster to a directly attached tape device.
- **Configuring full cluster backups to an IBM Spectrum Protect server**
You can configure IBM Spectrum Protect to back up all volumes in a NetApp cluster to an IBM Spectrum Protect server, which stores the data in a storage pool. Even if you move volumes within the cluster, backup operations continue and no reconfiguration is required.
- **Configuring partial cluster backups to an IBM Spectrum Protect server**
You can configure IBM Spectrum Protect to complete a partial backup of a NetApp cluster. This method is useful when multiple organizations own data in the cluster. Each organization can manage backup operations for its data.
- **Reconfiguring IBM Spectrum Protect to optimize clustered backups**
If you configured IBM Spectrum Protect to back up NetApp clusters by using node-scoped NDMP, you can reconfigure IBM Spectrum Protect to use NDMP Cluster Aware Backup (CAB). In this way, you can optimize backup operations for NetApp clusters.

Configuring full cluster backups to directly attached tape devices

You can configure IBM Spectrum Protect™ to back up all volumes in a NetApp cluster to a directly attached tape device.

Before you begin

For an overview of NDMP functionality in IBM Spectrum Protect and NetApp file servers, see technote 7046965. This technote also lists system requirements.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is installed on your NetApp file server, use the following procedure. After you configure your NetApp file server to work with IBM Spectrum Protect, you can use the NetApp Cluster Aware Backup (CAB) extension to back up all volumes.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is not installed on your NetApp file server, back up data by following the instructions in Configuring IBM Spectrum Protect for NDMP operations in a nonclustered environment.

About this task

The preferred method is to back up the full cluster by using a node and data mover that are associated with the clusterwide network. In this way, you ensure that the backup data is owned by a single IBM Spectrum Protect node. Even if you move volumes within the cluster, backup operations continue and no reconfiguration is required.

Procedure

To configure full cluster backup operations to a directly attached tape device, complete the following steps:

1. Verify that IBM Spectrum Protect Extended Edition is installed, and that the license is registered. If the license is not registered, issue the following IBM Spectrum Protect command:

```
register license file=tsmee.lic
```

2. Obtain cluster administrator privileges for the NetApp file server. This step is required to access the cluster console.
3. On the NetApp file server, enable the use of NDMP by following the instructions in the *Clustered Data ONTAP® Data Protection Tape Backup and Recovery Guide*. Complete the following steps:
 - a. Enable SVM-scoped NDMP backup operations at the cluster level. In this way, you disable node-scoped NDMP backup operations on the NAS file server. Ensure that the node-scoped-ndmp option on the NAS file server is set to OFF.
 - b. Create a backup user ID for NDMP operations.
 - c. Configure a network interface for NDMP control connections at the cluster level.
4. Register the IBM Spectrum Protect node that will own all backup data for the cluster. On the IBM Spectrum Protect server, issue the REGISTER NODE command:

```
register node node_name password domain=nas_domain type=nas
```

where *node_name* specifies the node name, *password* specifies the node password, and *nas_domain* specifies the domain of the node. Assign the node to a domain that has a policy to back up data to an appropriate storage pool.

5. Determine the IP address of the NetApp cluster management interface on the NAS file server. The interface provides access to the entire cluster. On the NAS file server, issue the following Data ONTAP operating system command:

```
network interface show -role cluster-mgmt
```

The IP address that is shown in the command output is required when you specify the HLADDRESS parameter in step 6.

6. Define a data mover for the IBM Spectrum Protect node that will own the backup data. On the IBM Spectrum Protect server, issue the DEFINE DATAMOVER command on one line:

```
define datamover data_mover_name type=nascluster  
hladdress=cluster_management_interface lladdress=port  
USER=user_name password=password dataformat=netappdump
```

where *cluster_management_interface* is the value that you obtained in step 5 and *data_mover_name* is the node name that you registered in step 4. For information about specifying the other parameters, see DEFINE DATAMOVER (Define a data mover).

Tip: After you define the data mover, additional data movers are defined automatically for each node in the cluster. The name of each data mover matches the name of the physical node in the cluster. You will use these data movers when you define paths to tape drives in step 3 of Configuring tape devices for full cluster backups.

What to do next

To configure the tape device for the full cluster backup, follow the instructions in *Configuring tape devices for full cluster backups*.

- *Configuring tape devices for full cluster backups*

If you plan to back up all volumes in a NetApp cluster to a directly attached tape device, you must configure the tape device.

Related reference:

REGISTER NODE (Register a node)

Configuring full cluster backups to an IBM Spectrum Protect server

You can configure IBM Spectrum Protect™ to back up all volumes in a NetApp cluster to an IBM Spectrum Protect server, which stores the data in a storage pool. Even if you move volumes within the cluster, backup operations continue and no reconfiguration is required.

Before you begin

For an overview of NDMP functionality in IBM Spectrum Protect and NetApp file servers, see technote 7046965. This technote also lists system requirements.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is installed on your NetApp file server, use the following procedure. After you configure your NetApp file server to work with IBM Spectrum Protect, you can use the NetApp Cluster Aware Backup (CAB) extension to back up all volumes in the cluster. All backed-up data will be owned by a single IBM Spectrum Protect node.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is not installed on your NetApp file server, back up data by following the instructions in *Configuring IBM Spectrum Protect for NDMP operations in a nonclustered environment*.

Procedure

1. Verify that IBM Spectrum Protect Extended Edition is installed, and that the license is registered. If the license is not registered, issue the following IBM Spectrum Protect command:

```
register license file=tsmee.lic
```

2. Obtain cluster administrator privileges for the NetApp file server. This step is required to access the cluster console.
3. Enable the use of NDMP by following the instructions in the *Clustered Data ONTAP® Data Protection Tape Backup and Recovery Guide*. Complete the following steps:
 - a. Enable the NetApp SVM to control NDMP backup operations at the cluster level.
 - b. Create a backup user ID for NDMP operations.
 - c. Configure a network interface for NDMP control connections at the cluster level.
4. Register the IBM Spectrum Protect node that will own all backup data for the cluster. On the IBM Spectrum Protect server, issue the REGISTER NODE command:

```
register node node_name password domain=nas_domain type=nas
```

where *node_name* specifies the node name, *password* specifies the node password, and *nas_domain* specifies the domain of the node.

5. Determine the numerical IP address or the domain name that is used to access the NAS file server. The interface provides access to the entire cluster. On the NAS file server, issue the following Data ONTAP operating system command:

```
network interface show -role cluster-mgmt
```

The IP address in the output is required when you specify a value for the HLADDRESS parameter in step 6.

6. Define a data mover for the node by issuing the DEFINE DATAMOVER command and specifying TYPE=NASCLUSTER. On the IBM Spectrum Protect server, issue the following command on one line:

```
define datamover data_mover_name type=nascluster  
hladdress=cluster_management_interface lladdress=port  
USER=user_name password=password dataformat=netappdump
```

where *cluster_management_interface* is the value that you obtained in step 5 and *data_mover_name* is the node name that you registered in step 4. For information about specifying the other parameters, see DEFINE DATAMOVER (Define a data mover).

7. Configure an IBM Spectrum Protect policy for managing NAS image backups. Follow the instructions in Configuring an IBM Spectrum Protect policy for NDMP operations.
8. Update the cluster node that you registered in step 4 to the domain that was configured in step 7. On the IBM Spectrum Protect server, issue the UPDATE NODE command:

```
update node node_name domain=domain_name
```

9. Optional: Identify the volumes in the cluster and schedule backups for the volumes:
 - a. On the NAS file server, identify the volumes in the cluster by issuing the following Data ONTAP command:

```
volume show
```

- b. Schedule backup operations by following the instructions in Scheduling NDMP operations.

What to do next

The following tasks are optional:

- To verify that volumes in the NetApp cluster are backed up, complete the following steps:
 1. On the Operations Center menu bar, click Clients.
 2. Double-click a NAS device client and click Volumes.
 3. To determine when the last full volume backup was completed, review the information in the Last Full column. To determine when the most recent differential backup was completed, review the information in the Last Differential column.
- To set up copy storage pools for added data protection, configure the tape-to-tape function to back up data. For instructions, see Backing up data with the tape-to-tape function.

Related reference:

REGISTER NODE (Register a node)

Configuring partial cluster backups to an IBM Spectrum Protect server

You can configure IBM Spectrum Protect™ to complete a partial backup of a NetApp cluster. This method is useful when multiple organizations own data in the cluster. Each organization can manage backup operations for its data.

Before you begin

For an overview of NDMP functionality in IBM Spectrum Protect and NetApp file servers, see technote 7046965. This technote also lists system requirements.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is installed on your NetApp file server, use the following procedure. After you configure your NetApp file server to work with IBM Spectrum Protect, you can use the NetApp Cluster Aware Backup (CAB) extension to back up a partial cluster. When you configure a partial cluster backup, you determine the scope of the backup by specifying a virtual server, the NetApp storage virtual machine (SVM). The SVM provides access to part of a cluster.

If the NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, operating system is not installed on your NetApp file server, back up data by following the instructions in Configuring IBM Spectrum Protect for NDMP operations in a nonclustered environment.

Procedure

1. Verify that IBM Spectrum Protect Extended Edition is installed, and that the license is registered. If the license is not registered, issue the following IBM Spectrum Protect command:

```
register license file=tsmee.lic
```

2. Obtain cluster administrator privileges for the NetApp file server. This step is required to access the cluster console.
3. On the NetApp file server, enable the use of NDMP by following the instructions in the *Clustered Data ONTAP® Data Protection Tape Backup and Recovery Guide*. Complete the following steps:
 - a. Enable the NetApp SVM to control NDMP backup operations.
 - b. Create a backup user ID for NDMP operations.
 - c. Configure a network interface for NDMP control connections at the SVM level.
4. Register the IBM Spectrum Protect node that will own the backed-up data. On the IBM Spectrum Protect server, issue the REGISTER NODE command:

```
register node node_name password domain=nas_domain type=nas
```

where *node_name* specifies the node name, *password* specifies the node password, and *nas_domain* specifies the domain of the node.

5. Determine the numerical IP address or the domain name of the cluster interface that is used by the SVM. To determine the value, on the NAS file server, issue the following ONTAP operating system command:

```
network interface show -vserver vserver_name -role data
```

where *vserver_name* specifies the name of the SVM. This value that you obtain is required in step 6.

6. Define an associated data mover for the IBM Spectrum Protect node by issuing the DEFINE DATAMOVER command and specifying TYPE=NASVSERVER. On the IBM Spectrum Protect server, issue the following command on one line:

```
define datamover data_mover_name type=nasvserver  
hladdress=svm_data_interface lladdress=port  
USER=user_name password=password dataformat=netappdump
```

where *svm_data_interface* is the value that you obtained in step 5 and *data_mover_name* is the name of the node that you registered in step 4.

For information about specifying the other parameters, see DEFINE DATAMOVER (Define a data mover).

7. Configure an IBM Spectrum Protect policy for managing NAS image backups. Follow the instructions in Configuring an IBM Spectrum Protect policy for NDMP operations.
8. Update the node that you registered in step 4 to the domain that you configured in step 7. On the IBM Spectrum Protect server, issue the UPDATE NODE command:

```
update node node_name domain=domain_name
```

9. Optional: Identify the volumes in the cluster and schedule backup operations. Take the following steps:
 - a. On the NAS file server, identify the volumes in the cluster by issuing the following Data ONTAP command:

```
volume show -vserver vserver_name
```

where *vserver_name* specifies the name of the SVM.

- b. Schedule backup operations by following the instructions in Scheduling NDMP operations.

What to do next

To verify that volumes in the NetApp cluster are backed up, complete the following steps:

1. On the Operations Center menu bar, click Clients.
2. Double-click a NAS device client and click Volumes.
3. To determine when the last full volume backup was completed, review the information in the Last Full column. To determine when the most recent differential backup was completed, review the information in the Last Differential column.

Related reference:

REGISTER NODE (Register a node)

Reconfiguring IBM Spectrum Protect to optimize clustered backups

If you configured IBM Spectrum Protect™ to back up NetApp clusters by using node-scoped NDMP, you can reconfigure IBM Spectrum Protect to use NDMP Cluster Aware Backup (CAB). In this way, you can optimize backup operations for NetApp clusters.

Before you begin

For an overview of NDMP functionality in IBM Spectrum Protect and NetApp file servers, see technote 7046965. This technote also lists system requirements.

About this task

When you reconfigure IBM Spectrum Protect to use CAB, you can optimize backup operations in the following ways:

- You can configure IBM Spectrum Protect to back up all volumes in a NetApp cluster to a directly attached tape device or to an IBM Spectrum Protect server. In both cases, the data is owned by a single IBM Spectrum Protect node. Even if you move

volumes within the cluster, backup operations continue and no reconfiguration is required.

- You can complete a partial backup of a NetApp cluster to an IBM Spectrum Protect server. This method is useful when multiple organizations own data in the cluster. Each organization can manage backup operations for its data. You set the scope of a partial backup by specifying a NetApp storage virtual machine (SVM), which provides access to part of a cluster.

To reconfigure IBM Spectrum Protect to use CAB, you must define a new IBM Spectrum Protect node and a new data mover.

Procedure

1. Verify that NetApp Clustered Data ONTAP 8.2 or later, or 9.1 or later, is installed on the NetApp file server.
2. Enable the use of NDMP by following the instructions in the *Clustered Data ONTAP® Data Protection Tape Backup and Recovery Guide*. Take one of the following actions:

For a full cluster backup

Complete the following steps:

- a. Enable SVM-scoped NDMP backup operations at the cluster level. In this way, you disable node-scoped NDMP backup operations on the NAS file server. Ensure that the node-scoped-ndmp option on the NAS file server is set to OFF.
- b. Create a backup user ID for NDMP operations.
- c. Configure a network interface for NDMP control connections at the cluster level.

For a partial cluster backup

Complete the following steps:

- a. Enable SVM-scoped NDMP to control NDMP backup operations.
- b. Create a backup user ID for NDMP operations.
- c. Configure a network interface for NDMP control connections at the SVM level.

3. Register the IBM Spectrum Protect node that will own the backup data. On the IBM Spectrum Protect server, issue the REGISTER NODE command:

```
register node node_name password password domain=nas_domain type=nas
```

where *node_name* specifies the node name, *password* specifies the node password, and *nas_domain* specifies the domain of the node.

4. If you plan to back up a full cluster, determine the IP address of the NetApp cluster management interface on the NAS file server. The interface provides access to the entire cluster. On the NAS file server, issue the following Data ONTAP operating system command:

```
network interface show -role cluster-mgmt
```

The IP address in the output is required when you specify the HLADDRESS parameter in step 6.

5. If you plan to back up a partial cluster, determine the numerical IP address or the domain name of the cluster interface that is used by the SVM. To determine the value, issue the following Data ONTAP operating system command on the NAS file server:

```
network interface show -vserver vserver_name -role data
```

where *vserver_name* specifies the name of the SVM. The value that you obtain is required in step 6.

6. Define a data mover for the IBM Spectrum Protect node. Take one of the following actions:

For a full cluster backup

Define a data mover for the IBM Spectrum Protect node that will own the backup data. On the IBM Spectrum Protect server, issue the DEFINE DATAMOVER command on one line:

```
define datamover data_mover_name type=nascluster  
hladdress=cluster_management_interface lladdress=port  
USER=user_name password=password dataformat=netappdump
```

where *cluster_management_interface* is the value that you obtained in step 4 and *data_mover_name* is the node name that you registered in step 3.

Tip: After you define the data mover, additional data movers are defined automatically for each node in the cluster. The name of each data mover matches the name of the physical node in the cluster. You will use these data movers when you define paths to tape drives that are attached to the cluster.

For a partial cluster backup

Define a data mover for the node by issuing the DEFINE DATAMOVER command and specifying TYPE=NASVSERVER. On the IBM Spectrum Protect server, issue the following command on one line:

```
define datamover data_mover_name type=nasvserver
hladdress=svm_data_interface lladdress=port
USER=user_name password=password dataformat=netappdump
```

where *svm_data_interface* is the value that you obtained in step 5 and *data_mover_name* is the node name that you registered in step 3.

For information about specifying the other parameters on the DEFINE DATAMOVER command, see DEFINE DATAMOVER (Define a data mover).

7. To back up data to a directly attached tape device, for each tape drive that is attached to the cluster, identify the device name and the physical node to which the drive is attached:
 - a. On the NAS file server, issue the following Data ONTAP command:

```
storage tape show-tape-drive
```
 - b. Review the output to find the serial number of the tape drive, and the node of the cluster to which the drive is attached. The same stanza includes the device name, for example, *st1*, *st2*, or *st3*.
8. To configure a full cluster backup to a directly attached tape device, follow the instructions in Configuring tape devices for full cluster backups.
9. To configure a full or partial cluster backup to an IBM Spectrum Protect server, configure a policy for managing NAS image backups. Follow the instructions in Configuring an IBM Spectrum Protect policy for NDMP operations.
10. Disable scheduled backup operations for all nodes that were previously used to back up the NetApp cluster.
11. Identify the volumes in the cluster and optionally schedule backup operations for the volumes. Take one of the following actions:

For a full cluster backup

- a. On the NAS file server, identify the volumes in the cluster by using the following Data ONTAP command:

```
volume show
```

- b. Run a full backup of the entire cluster.
- c. Optional: To schedule backup operations, follow the instructions in Scheduling NDMP operations.

For a partial cluster backup

- a. On the NAS file server, identify the volumes in the cluster by using the following Data ONTAP command:

```
volume show -vserver vserver_name
```

where *vserver_name* specifies the name of the SVM.

- b. Run a full backup of the partial cluster.
- c. Optional: To schedule backup operations, follow the instructions in Scheduling NDMP operations.

What to do next

To verify that volumes in the NetApp cluster are backed up, complete the following steps:

1. On the Operations Center menu bar, click Clients.
2. Double-click a NAS device client and click Volumes.
3. To determine when the last full volume backup was completed, review the information in the Last Full column. To determine when the most recent differential backup was completed, review the information in the Last Differential column.

Related reference:

DEFINE DATAMOVER (Define a data mover)
DEFINE PATH (Define a path when the destination is a drive)
REGISTER NODE (Register a node)

Backing up and restoring NAS file servers using NDMP

After you configure IBM Spectrum Protect™ for NDMP operations, you are ready to begin using NDMP.

Procedure

Use either a client interface or an administrative interface to perform a file system image backup. For example, to use the Windows backup-archive client interface to back up a file system that is named */vol/vol1* on a NAS file server that is named *NAS1*, issue the following command:

```
dsmc backup nas -nasnodename=nas1 {/vol/vol1}
```

For more information about the command, see Backup Image.

Tip: Whenever you use the client interface, you are asked to authenticate yourself as an IBM Spectrum Protect administrator before the operation can begin. The administrator ID must have at least client owner authority for the NAS node.

You can complete the same backup operation with a server interface. For example, from the administrative command-line client, back up the file system that is named /vol/vol1 on a NAS file server that is named NAS1, by issuing the following command:

```
backup node nas1 /vol/vol1
```

Restriction: The BACKUP NAS and BACKUP NODE commands do not include snapshots. To back up snapshots, see Backing up and restoring with snapshots.

You can restore the image by using either interface. Backups are identical whether they are backed up using a client interface or a server interface. For example, suppose that you want to restore the image that is backed up in the previous examples. For this example, the file system that is named /vol/vol1 is being restored to /vol/vol2. Restore the file system with the following command, issued from a Windows backup-archive client interface:

```
dsmc restore nas -nasnodename=nas1 {/vol/vol1} {/vol/vol2}
```

You can choose to restore the file system by using a server interface. For example, to restore the file system name /vol/vol1 to file system /vol/vol2, for a NAS file server that is named NAS1, enter the following command:

```
restore node nas1 /vol/vol1 /vol/vol2
```

You can restore data from one NAS vendor system to another NAS vendor system when you use the NDMPDUMP data format. However, you must either verify compatibility between systems or maintain a separate storage pool for each NAS vendor.

- NAS file servers: backups to a single IBM Spectrum Protect server
If you have several NAS file servers in different locations, you might prefer to send the backup data to a single IBM Spectrum Protect server rather than attaching a tape library to each NAS device.
- Backing up NDMP file servers to an IBM Spectrum Protect server
You can back up data to a single IBM Spectrum Protect server rather than attaching a tape library to each NAS device.

NAS file servers: backups to a single IBM Spectrum Protect server

If you have several NAS file servers in different locations, you might prefer to send the backup data to a single IBM Spectrum Protect™ server rather than attaching a tape library to each NAS device.

When you store NAS backup data in the IBM Spectrum Protect server's storage hierarchy, you can apply IBM Spectrum Protect back-end data management functions. In this way, you can take advantage of migration, reclamation, disaster recovery, and other features.

To back up a NAS device to an IBM Spectrum Protect native storage pool, set the destination storage pool in the copy group to point to the wanted native storage pool. The destination storage pool provides the information about the library and drives that are used for backup and restore. You must ensure that there is sufficient space in your target storage pool to contain the NAS data, which can be backed up to sequential, disk, or file type devices. Defining a separate device class is not necessary.

If you are creating a table of contents, a management class must be specified with the TOCDESTINATION parameter in the DEFINE and UPDATE COPYGROUP commands. When you back up a NAS file server to IBM Spectrum Protect native pools, the TOCDESTINATION can be the same as the destination of the data that is backed up by using NDMP.

Firewall considerations are more stringent than they are for filer-to-attached-library because communications can be initiated by either the IBM Spectrum Protect server or the NAS file server. NDMP tape servers run as threads within the IBM Spectrum Protect server and the tape server accepts connections on port of 10001. This port number can be changed through the following option in the IBM Spectrum Protect server options file: NDMPPORTRANGE port-number-low, port-number-high.

During NDMP filer-to-server backup operations, you can use the NDMPREFDATAINTERFACE option to specify which network interface the IBM Spectrum Protect server uses to receive backup data. The value for this option is a host name or IPV4 address that is associated with one of the active network interfaces of the system on which the IBM Spectrum Protect server is running. This interface must be IPV4 enabled.

Before you use this option, verify that your NAS device supports NDMP operations that use a different network interface for NDMP control and NDMP data connections. NDMP control connections are used by IBM Spectrum Protect to authenticate with an NDMP server and monitor an NDMP operation while NDMP data connections are used to transmit and receive back up data during NDMP operations. You must still configure your NAS device to route backup and restore data to the appropriate network interface.

When enabled, the NDMPREFDATAINTERFACE option affects all subsequent NDMP filer-to-server operations. It does not affect NDMP control connections because they use the system's default network interface. You can update this server option without stopping and restarting the server by using the SETOPT command.

NetApp file servers provide an NDMP option (ndmpd.preferred_interface) to change the interface that is used for NDMP data connections. For more information, see the documentation for your NAS device.

For instructions about completing NDMP filer-to-server backup operations, see Backing up NDMP file servers to an IBM Spectrum Protect server.

For information about server options, see Server options.

Backing up NDMP file servers to an IBM Spectrum Protect server

You can back up data to a single IBM Spectrum Protect™ server rather than attaching a tape library to each NAS device.

Procedure

To back up a server on a NAS file system, complete the following steps:

1. Select an existing storage pool or set up a storage pool for the NAS data by issuing the following command:

```
define stgpool naspool disk
```

2. Define volumes to add to the storage pool. For example, define a volume that is named naspool_volAB:

```
define volume naspool /usr/storage/naspool_volAB formatsize=100
```

3. Set the copy destination to the storage pool defined previously and activate the associated policy set.

```
update copygroup standard standard standard destination=naspool  
tocdestination=naspool  
activate policyset standard standard
```

The destination for NAS data is determined by the destination in the copy group. The storage size estimate for NAS differential backups uses the occupancy of the file space, the same value that is used for a full backup. You can use this size estimate as one of the considerations in choosing a storage pool. One of the attributes of a storage pool is the MAXSIZE value, which indicates that data is sent to the NEXT storage pool when the MAXSIZE value is exceeded by the estimated size. Because NAS differential backups to IBM Spectrum Protect native storage pools use the base file space occupancy size as a storage size estimate, differential backups end up in the same storage pool as the full backup. Depending on collocation settings, differential backups might end up on the same media as the full backup.

4. Set up a node and data mover for the NAS device. The data format signifies that the backup images created by this NAS device are a dump type of backup image in a NetApp specific format.

```
register node nas1 nas1 type=nas domain=standard  
define datamover nas1 type=nas hla=nas1 user=root  
password=***** dataformat=netappdump
```

The NAS device is now ready to be backed up to an IBM Spectrum Protect server storage pool. Paths can be defined to local drives, but the destination that is specified by the management class determines the target location for this backup operation.

5. Back up the NAS device to the IBM Spectrum Protect storage pool by issuing the following command:

```
backup node nas1 /vol/vol0
```

6. Restore a NAS device from the IBM Spectrum Protect storage pool by issuing the following command:

```
restore node nas1 /vol/vol0
```

File-level backup and restore for NDMP operations

When you back up data by using NDMP, you can specify that the IBM Spectrum Protect™ server collects and stores file-level information in a table of contents (TOC).

If you specify this option at the time of backup, you can later display the TOC of the backup image. By using the backup-archive web client Version 8.1.1 or earlier, you can select individual files or directories to restore data directly from the generated backup

images.

Restriction: If you installed the backup-archive client V8.1.1 or earlier, you can use the web client interface for file-level restore operations. If you installed the backup-archive client V8.1.2 or later, you cannot use the web client interface for file-level restore operations.

Collecting file-level information requires extra processing time, network resources, storage pool space, temporary database space, and possibly additional storage device interaction. For instructions about configuring storage devices, see [Configuring storage devices](#). You must consider dedicating more space in the IBM Spectrum Protect server database. You must set up policy so that the IBM Spectrum Protect server stores the TOC in a different storage pool from the one where the backup image is stored. The TOC is treated like any other object in that storage pool.

You can also complete a backup operation by using NDMP without collecting file-level restore information.

To allow creation of a TOC for a backup by using NDMP, you must define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. You cannot specify a copy storage pool or an active-data pool as the destination. The storage pool that you specify for the TOC destination must have a data format of either NATIVE or NONBLOCK, so it cannot be the tape storage pool that is used for the backup image.

If you choose to collect file-level information, specify the TOC parameter in the BACKUP NODE server command. Or, if you initiate your backup by using the client, you can specify the TOC option in the client options file, client option set, or client command line. You can specify NO, PREFERRED, or YES. When you specify PREFERRED or YES, the IBM Spectrum Protect server stores file information for a single NDMP-controlled backup in a TOC. The TOC is placed into a storage pool. After that, the IBM Spectrum Protect server can access the TOC so that file and directory information can be queried by the server or client. Use of the TOC parameter allows a TOC to be generated for some images and not others, without requiring different management classes for the images.

For more information about the BACKUP NODE command, see [BACKUP NODE \(Back up a NAS node\)](#).

To avoid mount delays and ensure sufficient space, use random access storage pools (DISK device class) as the destination for the TOC. For sequential access storage pools, no labeling or other preparation of volumes is necessary if scratch volumes are allowed.

For more information, see [Managing tables of contents](#).

- Interfaces for file-level restore operations
To restore individual files and directories, you can use one of the following interfaces: the backup-archive web client at Version 8.1.1 or earlier, or the server interface.
- International characters for NetApp file servers
All systems that create or access data on a particular NAS file server volume must do so in a manner compatible with the volume language setting.
- File-level restore operations from a directory-level backup image
File-level restore operations are supported for directory-level backup images.

Interfaces for file-level restore operations

To restore individual files and directories, you can use one of the following interfaces: the backup-archive web client at Version 8.1.1 or earlier, or the server interface.

Restriction: If you installed the backup-archive client V8.1.1 or earlier, you can use the web client interface for file-level restore operations. If you installed the backup-archive client V8.1.2 or later, you cannot use the web client interface for file-level restore operations.

Guidelines for using the backup-archive web client V8.1.1 or earlier:

To restore files and directories, a TOC must exist. The web client must be on a Windows system. The IBM Spectrum Protect™ server accesses the TOC from the storage pool and loads TOC information into a temporary database table. Then, you can use the web client to examine directories and files that are contained in one or more file system images. You can select individual files or directories to restore data directly from the generated backup images.

Guidelines for using the server interface:

- If you have a TOC, you can display the files in the backup by using the QUERY TOC command. When you run the RESTORE NODE command, specify one or more files from the output on the FILELIST parameter.
- If you did not create a TOC, the contents of the backup image cannot be displayed. You can restore individual files, directories, or both if you know the name of the file or directory, and in which image the backup is located. Use the RESTORE NODE command with the FILELIST parameter.

International characters for NetApp file servers

All systems that create or access data on a particular NAS file server volume must do so in a manner compatible with the volume language setting.

You must install Data ONTAP 6.4.1 or later, if it is available, on your NetApp NAS file server to garner full support of international characters in the names of files and directories.

If your level of Data ONTAP is earlier than 6.4.1, you must have one of the following two configurations to collect and restore file-level information. Results with configurations other than the two listed are unpredictable. The IBM Spectrum Protect™ server issues a warning message (ANR4946W) during backup operations. The message indicates that the character encoding of NDMP file history messages is unknown, and UTF-8 is assumed to build a table of contents. It is safe to ignore this message only for the following two configurations.

- Your data has directory and file names that contain only English (7-bit ASCII) characters.
- Your data has directory and file names that contain non-English-language characters and the volume language is set to the UTF-8 version of the proper locale (for example, `de.UTF-8` for German).

If your level of Data ONTAP is 6.4.1 or later, you must have one of the following three configurations to collect and restore file-level information. Results with configurations other than the three listed are unpredictable.

- Your data has directory and file names that contain only English (7-bit ASCII) characters and the volume language is either not set or is set to one of the following values:
 - `C` (POSIX)
 - `en`
 - `en_US`
 - `en.UTF-8`
 - `en_US.UTF-8`
- Your data has directory and file names that contain non-English-language characters, and the volume language is set to the proper locale (for example, `de.UTF-8` or `de` for German).
Tip: Using the UTF-8 version of the volume language setting is more efficient in terms of IBM Spectrum Protect server processing and table of contents storage space.
- You use CIFS only to create and access your data.

File-level restore operations from a directory-level backup image

File-level restore operations are supported for directory-level backup images.

As with a NAS file system backup, a table of contents (TOC) is created during a directory-level backup. You can browse the files in the image by using the web client Version 8.1.1 or earlier. By default, the files are restored to the original location. However, during a file-level restore from a directory-level backup, you can either select a different file system or another virtual file space name as a destination.

Restriction: If you installed the backup-archive client V8.1.1 or earlier, you can use the web client interface for file-level restore operations. If you installed the backup-archive client V8.1.2 or later, you cannot use the web client interface for file-level restore operations.

For a TOC of a directory-level backup image, the path names for all files are relative to the directory that is specified in the virtual file space definition, not the root of the file system.

Directory-level backup and restore operations

If you have a large NAS file system, initiating a backup at a directory level reduces backup and restore times and provides more flexibility in configuring NAS backups. By defining virtual file spaces, a file system backup can be partitioned among several NDMP backup operations and multiple tape drives. You can also use different backup schedules to back up subtrees of a file system.

The virtual file space name cannot be identical to any file system on the NAS node. If a file system is created on the NAS device with the same name as a virtual file system, a name conflict occurs on the IBM Spectrum Protect™ server when the new file space is backed up. For instructions about issuing commands for mapping virtual file spaces, see `DEFINE VIRTUALFSMAPPING` (Define a virtual file space mapping).

Restriction: Virtual file space mappings are supported only for NAS nodes.

- Directory-level backup and restore for NDMP operations
The DEFINE VIRTUALFSMAPPING command maps a directory path of a NAS file server to a virtual file space name on the IBM Spectrum Protect server. After a mapping is defined, you can conduct NAS operations such as BACKUP NODE and RESTORE NODE by using the virtual file space names as if they were actual NAS file spaces.
- Backing up and restoring with snapshots
NDMP directory-level backup operations give you the ability to back up user-created snapshots of a NAS file system. Those snapshots are then stored as subdirectories. The snapshots can be taken at any time, and the backup to tape can be deferred to a more convenient time.

Directory-level backup and restore for NDMP operations

The DEFINE VIRTUALFSMAPPING command maps a directory path of a NAS file server to a virtual file space name on the IBM Spectrum Protect™ server. After a mapping is defined, you can conduct NAS operations such as BACKUP NODE and RESTORE NODE by using the virtual file space names as if they were actual NAS file spaces.

To start a backup of the directory, issue the BACKUP NODE command and specify the virtual file space name instead of a file space name. To restore the directory subtree to the original location, run the RESTORE NODE command and specify the virtual file space name.

Virtual file space definitions can also be specified as the destination in a RESTORE NODE command. In this way, you can restore backup images (either file system or directory) to a directory on any file system of the NAS device.

Backing up and restoring with snapshots

NDMP directory-level backup operations give you the ability to back up user-created snapshots of a NAS file system. Those snapshots are then stored as subdirectories. The snapshots can be taken at any time, and the backup to tape can be deferred to a more convenient time.

Procedure

For example, to back up a snapshot that is created for a NetApp file system, complete the following steps:

1. On the console for the NAS device, issue the command to create the snapshot. SNAP CREATE is the command for a NetApp device.

```
snap create vol2 february17
```

This example creates a snapshot that is named FEBRUARY 17 of the /vol/vol2 file system. The physical location for the snapshot data is in the directory /vol/vol2/.snapshot/february17. The stored location for snapshot data depends on the NAS vendor implementation. For NetApp, the SNAP LIST command can be used to display all snapshots for a file system.

2. Define a virtual file space mapping definition on the IBM Spectrum Protect™ server for the snapshot data that is created in the previous step.

```
define virtualfsmapping nas1 /feb17snapshot /vol/vol2 /.snapshot/february17
```

This example creates a virtual file space mapping definition named /feb17snapshot.

3. Back up the virtual file space mapping.

```
backup node nas1 /feb17snapshot mode=full toc=yes
```

4. After the backup is created, you can either restore the entire snapshot image or restore an individual file. Before you restore the data, you can create a virtual file space mapping name for the target directory. You can select any file system name as a target. The target location in this example is the directory /feb17snaprestore on the file system /vol/vol1.

```
define virtualfsmapping nas1 /feb17snaprestore /vol/vol1 /feb17snaprestore
```

5. Restore the snapshot backup image.

```
restore node nas1 /feb17snapshot /feb17snaprestore
```

This example restores a copy of the /vol/vol2 file system to the directory /vol/vol1/feb17snaprestore in the same state as when the snapshot was created in the first step.

Backup and restore operations by using the NetApp SnapMirror to Tape feature

You can back up large NetApp file systems by using the NetApp SnapMirror to Tape feature (also known as SMTape). Using a block-level copy of data for backup, the SnapMirror to Tape method is faster than a traditional NDMP full backup and can be used when NDMP full backups are impractical.

Use the NDMP SnapMirror to Tape feature as a disaster recovery option for copying large NetApp file systems to auxiliary storage. For most NetApp file systems, use the standard NDMP full or differential backup method.

By specifying a parameter on the BACKUP NODE and RESTORE NODE commands, you can back up and restore file systems by using SnapMirror to Tape. There are several limitations and restrictions on how SnapMirror images can be used. Consider the following guidelines before you use it as a backup method:

- If you installed NetApp ONTAP 8.2 or later, you must define a data mover of type NASCLUSTER or NASVSERVER for SnapMirror to Tape operations.
- You cannot initiate a SnapMirror to Tape backup or restore operation from the IBM Spectrum Protect™ Operations Center, web client, or command-line client.
- You cannot perform differential backups of SnapMirror images.
- You cannot perform a directory-level backup by using SnapMirror to Tape. Therefore, IBM Spectrum Protect does not permit SnapMirror to Tape backup operations on a server virtual file space.
- You cannot perform an NDMP file-level restore operation from SnapMirror to Tape images. Therefore, a table of contents is never created during SnapMirror to Tape image backups.
- At the start of a SnapMirror to Tape copy operation, the file server generates a snapshot of the file system. NetApp provides an NDMP environment variable to control whether this snapshot is removed at the end of the SnapMirror to Tape operation. IBM Spectrum Protect always sets this variable to remove the snapshot.
- After a SnapMirror to Tape image is retrieved and copied to a NetApp file system, the target file system remains configured as a SnapMirror partner. NetApp provides an NDMP environment variable to control whether this SnapMirror relationship should be broken. IBM Spectrum Protect always "breaks" the SnapMirror relationship during the retrieval. After the restore operation is complete, the target file system is in the same state as that of the original file system at the time of backup.

For more information about the SnapMirror to Tape feature, see BACKUP NODE (Back up a NAS node) and RESTORE NODE (Restore a NAS node).

NDMP backup operations using Celerra file server-integrated checkpoints

When the IBM Spectrum Protect™ server initiates an NDMP backup operation on a Celerra data mover, the backup of a large file system might take several hours to complete. Without Celerra integrated checkpoints, any changes that occur on the file system are written to the backup image.

As a result, the backup image includes changes that are made to the file system during the entire backup operation. The backup image is not a true point-in-time image of the file system.

If you are performing NDMP backup operations from Celerra file servers, upgrade the operating system of your data mover to Celerra file server version T5.5.25.1 or later. This version of the operating system allows enablement of integrated checkpoints for all NDMP backup operations from the Celerra Control Workstation. By enabling this feature, you ensure that the backup data represents true point-in-time images of the file system that is being backed up.

For instructions about enabling integrated checkpoints during all NDMP backup operations, see the Celerra file server documentation.

If your version of the Celerra file server operating system is earlier than version T5.5.25.1 and if you use NDMP to back up Celerra data movers, manually generate a snapshot of the file system by using Celerra's command-line checkpoint feature. Then, initiate an NDMP backup operation for the checkpoint file system rather than the original file system.

For instructions about creating and scheduling checkpoints from the Celerra control workstation, see the Celerra file server documentation.

Replicating NAS nodes

You can replicate a NAS node that uses NDMP for backup operations. Before you configure the replication operation, review the restrictions that apply.

About this task

Restrictions:

- The backup data must be in a storage pool with the NATIVE data format. You cannot replicate backup data in storage pools that have the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
- A differential backup can be replicated only if its full backup is replicated.

Procedure

1. Enable the NAS node for replication by issuing the UPDATE NODE command:

```
update node node_name replstate=enabled
```

where *node_name* specifies the name of the NAS node.

2. Replicate the node by issuing the REPLICATE NODE command:

```
replicate node node_name
```

where *node_name* specifies the name of the NAS node.

3. To ensure that the replicated data can be restored, define a data mover on the target server for the node by issuing the DEFINE DATAMOVER command:

```
define datamover node_name type=nas hladdress=hl_address lladdress=ll_address  
userid=user_id password=user_password dataformat=netappdump
```

where:

node_name

Specifies the name of the NAS node.

hl_address

Specifies either the numerical IP address or the domain name that is used to access the NAS file server.

ll_address

Specifies the TCP port number to access the NAS device for NDMP sessions.

user_id

Specifies the ID of a user who is authorized to initiate an NDMP session with the NAS file server.

user_password

Specifies the password of the user who is authorized to initiate an NDMP session with the NAS file server.

Results

The format of the backup data does not change during the replication process. If backup data is replicated, its associated table of contents is also replicated.

Data protection by using the NetApp SnapLock licensed feature

You can use the NetApp SnapLock licensed feature to meet strict regulatory requirements for archived data. When you enable the SnapLock feature, you can use IBM Spectrum Protect™ to set a retention date for files and to commit a file to a Write Once Read Many (WORM) state.

Data that is stored with a retention date cannot be deleted from the file system before the retention period expires. The SnapLock feature can be used by IBM Spectrum Protect servers only if the servers are enabled for data retention protection.

Data that is archived by data retention protection servers and stored to NetApp NAS file servers is stored as IBM Spectrum Protect FILE volumes. At the end of a write transaction, a retention date is set for the FILE volume, through the SnapLock interface. This date is calculated by using the RETVER and RETMIN parameters of the archive copy group that is used when you archive the data. By associating a retention date with the FILE volume, the FILE volume does not destroy or overwrite the data until the retention date passes. These FILE volumes are referred to as WORM FILE volumes. After a retention date is set, the WORM FILE volume

cannot be deleted until the retention date passes. IBM Spectrum Protect for Data Retention combined with WORM FILE volume reclamation ensures protection for the life of the data.

Storage pools can be managed either by threshold or by data retention period. The RECLAMATIONTYPE storage pool parameter indicates that a storage pool is managed based on a data retention period. When a traditional storage pool is queried with the FORMAT=DETAILED parameter, this output is displayed:

```
Reclamation Type: THRESHOLD
```

If an IBM Spectrum Protect server is enabled with data retention protection through IBM Spectrum Protect for Data Retention, and the server has access to a NetApp filer with the SnapLock licensed feature, you can define a storage pool with the RECLAMATIONTYPE parameter set to SNAPLOCK. This means that data that is created on volumes in this storage pool is managed by retention date. When a SnapLock storage pool is queried with the FORMAT=DETAILED parameter, the output indicates that the storage pools are managed by data retention period:

```
Reclamation Type: SNAPLOCK
```

For more information about the SnapLock filer, see the NetApp documentation *Data ONTAP Archive and Compliance Management Guide for 7-Mode*.

Attention: Do not use this feature to protect data with a retention period of less than three months.

- Reclamation and the SnapLock feature
To help ensure that data is always protected, set the NetApp default retention period to 30 days to match the default reclamation period of the WORM FILE volume. IBM Spectrum Protect reclaims any remaining data on a WORM FILE volume just before the retention date expiration.
- Retention periods
IBM Spectrum Protect policies manage the retention time for the WORM FILE volume. The retention of some files might exceed the retention time for the WORM FILE volume that they are stored on. You might need to move some files to another volume to ensure that the files are stored on WORM media.
- Configuration of the SnapLock feature for event-based retention
Data that is stored in SnapLock volumes that are managed by IBM Spectrum Protect for Data Retention and event-based retention can result in excessive reclamation, which causes performance degradation of the server.
- Continuous data protection with the SnapLock feature
If data is stored on a volume with the SnapLock feature enabled, and the data is moved or copied to a non-SnapLock volume, the data loses the unique hardware protection that is provided by NetApp WORM volumes.
- Setting up SnapLock volumes as IBM Spectrum Protect WORM FILE volumes
To meet strict requirements for archived data, enable the NetApp SnapLock feature.

Reclamation and the SnapLock feature

To help ensure that data is always protected, set the NetApp default retention period to 30 days to match the default reclamation period of the WORM FILE volume. IBM Spectrum Protect™ reclaims any remaining data on a WORM FILE volume just before the retention date expiration.

The reclamation of a WORM FILE volume to another WORM FILE volume before the retention date expiration helps to ensure that data is always protected by the SnapLock feature.

Because this protection is at an IBM Spectrum Protect volume level, the data on the volumes can be managed by IBM Spectrum Protect policy without consideration of where the data is stored. Data that is stored on WORM FILE volumes is protected both by data retention protection and by the retention period that is stored with the physical file on the SnapLock volume. If an IBM Spectrum Protect administrator issues a command to delete the data, the command fails. If someone attempts to delete the file by using a series of network file system calls, the SnapLock feature prevents the data from being deleted.

During reclamation processing, if the IBM Spectrum Protect server cannot move data from an expiring SnapLock volume to a new SnapLock volume, a warning message is issued.

Retention periods

IBM Spectrum Protect™ policies manage the retention time for the WORM FILE volume. The retention of some files might exceed the retention time for the WORM FILE volume that they are stored on. You might need to move some files to another volume to ensure that the files are stored on WORM media.

Some objects on the volume might need to be retained longer than other objects on the volume for the following reasons:

- The objects are bound to management classes with different retention times.
- The objects cannot be removed because of a deletion hold.
- The objects are waiting for an event to occur before expiration.
- The retention period for a copy group is increased, requiring a longer retention time than the time that is specified in the SnapLock feature when the WORM FILE volume was committed.

To manage a WORM FILE volume by retention time, you must issue the `DEFINE STGPOOL` command and specify `RECLAMATIONTYPE=SNAPLOCK`. In this way, you define a storage pool as a SnapLock storage pool. After that, you cannot update the `RECLAMATIONTYPE` parameter to a value of `THRESHOLD`. When you define a SnapLock storage pool, the system verifies that the specified directories are in the device class are SnapLock WORM volumes. When a file class is defined and storage pools are created with the reclamation type of `SNAPLOCK`, all volumes must be WORM volumes or the operation fails. If a device class is updated to contain extra directories and SnapLock storage pools are assigned to the device class, the same check is made to ensure that all directories are SnapLock WORM volumes.

Three retention periods are available for the NetApp SnapLock feature. The retention periods must be configured correctly so that the IBM Spectrum Protect server can properly manage WORM data that is stored in SnapLock volumes. The IBM Spectrum Protect server sets the retention period for data that is stored on NetApp SnapLock volumes based on the values in the copy group for the data that is archived. The NetApp file server must not conflict with the ability of the IBM Spectrum Protect server to set the retention period. The preferred method is to configure the following settings for retention periods in the NetApp file server:

- **Minimum Retention Period.** Set the higher value: either 30 days or the minimum number of days that is specified by any copy group (by using a NetApp SnapLock file server for WORM FILE storage) for the data retention period. The copy group is the one in use that stores data on NetApp SnapLock volumes.
- **Maximum Retention Period.** Leave the default value of 30 years. This retention period allows the IBM Spectrum Protect server to set the actual volume retention period based on the settings in the archive copy group.
- **Default Retention Period.** Set to 30 days. If you do not set this value and you do not set the maximum retention period, each volume's retention period is set to 30 years. If this occurs, the IBM Spectrum Protect server cannot manage expiration and reuse of NetApp SnapLock volumes. As a result, no volume can be reused for 30 years.

With the NetApp SnapLock retention periods set, IBM Spectrum Protect can manage the data in SnapLock storage pools with maximum efficiency. For each volume that is in a `SNAPLOCK` storage pool, an IBM Spectrum Protect reclamation period is created. The IBM Spectrum Protect reclamation period has a start date, `BEGIN RECLAIM PERIOD`, and an end date, `END RECLAIM PERIOD`. You can view these dates by issuing the `QUERY VOLUME` command with the `FORMAT=DETAILED` parameter on a SnapLock volume. The output is similar to this example:

```
Begin Reclaim Period: 09/05/2017
End Reclaim Period: 10/06/2017
```

When IBM Spectrum Protect archives files to a SnapLock volume, the server tracks the latest expiration date of those files, and the `BEGIN RECLAIM PERIOD` value is set to that latest expiration date. When more files are added to the SnapLock volume, the starting date is set to that later date if you have a file with a later expiration date than the one currently on the volume. The start date is set to the latest expiration date for any file on that volume. The expectation is that all files on that volume are already either expired, or are expiring on that day. On the following day, no valid data remains on that volume.

The `END RECLAIM PERIOD` is set to a month later than the `BEGIN RECLAIM PERIOD`. The retention date set in the NetApp file server for that volume is set to the `END RECLAIM PERIOD` date. The NetApp file server prevents deletion of that volume until the `END RECLAIM PERIOD` date is reached. This date is approximately a month after the data has expired in the IBM Spectrum Protect server. When the IBM Spectrum Protect server calculates an `END RECLAIM PERIOD` date for a volume, and the date is later than the current `END RECLAIM PERIOD`, the date is reset in the NetApp file server for that volume to the later date. Resetting the data to a later date guarantees that the IBM Spectrum Protect WORM FILE volume is not deleted until all data on the volume expires, or the data is moved to another SnapLock volume.

The IBM Spectrum Protect reclamation period is the amount of time between the begin date and the end date. During the reclamation period, the IBM Spectrum Protect server deletes volumes on which all the data is expired, or moves files that are not expired on expiring SnapLock volumes to new SnapLock volumes with new dates. This month is critical to how the server safely and efficiently manages the data on WORM FILE volumes. Data on a SnapLock volume typically expires by the time the beginning date arrives, and the volume must be empty. When the end date arrives, the volume can be safely deleted from the IBM Spectrum Protect inventory and the SnapLock file server.

However, some events might cause valid data to be on a SnapLock volume:

- Expiration processing in the IBM Spectrum Protect server for that volume might be delayed or is incomplete.

- The retention parameters on the copy group or associated management classes might be altered for a file after it was archived, and that file is not going to expire for some time.
- A deletion hold might be placed on one or more of the files on the volume.
- Reclamation processing is either disabled or is encountering errors by moving data to new SnapLock volumes on a SnapLock storage pool.
- A file is waiting for an event to occur before the IBM Spectrum Protect server can begin the expiration of the file.

When the beginning date arrives and files are not expired on a SnapLock volume, the files must be moved to a new SnapLock volume with a new begin and end date. However, if expiration processing is delayed on the IBM Spectrum Protect server, and those files expire when expiration processing on the IBM Spectrum Protect server runs, it is inefficient to move those files to a new SnapLock volume. To ensure that unnecessary data movement does not occur for files that are due to expire, movement of files on expiring SnapLock volumes will be delayed by a number of days after the BEGIN RECLAIM PERIOD date. Since the data is protected in the SnapLock file server until the END RECLAIM PERIOD date, there is no risk to the data in delaying this movement. This allows IBM Spectrum Protect expiration processing to finish. After that number of days, if valid data is on an expiring SnapLock volume, the data is moved to a new SnapLock volume, thus continuing the protection of the data.

Since the data was initially archived, there might be changes in the retention parameters for that data (for example, changes in the management class or copy pool parameters) or there might be a deletion hold on that data. However, the data on that volume is protected by SnapLock only until the END RECLAIM PERIOD date. Data that is not expired is moved to new SnapLock volumes during the IBM Spectrum Protect reclamation period. If errors occur when data is moved to a new SnapLock volume, a warning message is issued indicating that the data will soon be unprotected. If the error persists, issue a MOVE DATA command for the problem volume.

Attention: Do not disable reclamation processing on a SnapLock storage pool. After the processing is disabled, the IBM Spectrum Protect server has no way to issue warning messages that data will become unprotected. This situation can also occur if reclamation and migration are disabled for the entire server (for example, NOMIGRRECL set in the server options file). Ensure that your data is protected when you manage SnapLock storage pools.

Configuration of the SnapLock feature for event-based retention

Data that is stored in SnapLock volumes that are managed by IBM Spectrum Protect™ for Data Retention and event-based retention can result in excessive reclamation, which causes performance degradation of the server.

If data is managed by event-based retention, IBM Spectrum Protect initially sets the retention period to the greater of the RETVER and RETMIN values for the archive copy group. When the volume enters the reclamation period and data that remains on the volume is moved, the retention period for the target volume is set to the remaining retention period of the data, which is typically 0. The new volume then enters the reclamation period shortly after the volume receives the data, resulting in the reclamation of volumes that were just created.

You can avoid this situation by using the RETENTIONEXTENSION server option. This option allows the server to set or extend the retention date of a SnapLock volume. You can specify a value in the range 30 - 9999 days. The default is 365 days.

When you select volumes in a SnapLock storage pool for reclamation, the server verifies whether the volume is within the reclamation period:

- If the volume is not within the reclamation period, no action is taken. The volume is not reclaimed, and the retention date is unchanged.
- If the volume is within the reclamation period, the server verifies whether the percentage of reclaimable space on the volume is greater than the reclamation threshold of the storage pool or of the threshold percentage that is passed in on the THRESHOLD parameter of a RECLAIM STGPOOL command:
 - If the reclaimable space is greater than the threshold, the server reclaims the volume and sets the retention date of the target volume is set to the greater of these values:
 - The remaining retention time of the data plus 30 days for the reclamation period.
 - The RETENTIONEXTENSION value plus 30 days for the reclamation period.
 - If the reclaimable space is not greater than the threshold, the server resets the retention date of the volume by the amount that is specified in the RETENTIONEXTENSION option. The new retention period is calculated by adding the number of days that are specified to the current date.

In the following examples, the SnapLock volume, VolumeA, is in a storage pool whose reclamation threshold is set to 60%. The RETENTIONEXTENSION server option is set to 365 days. The retention period for VolumeA is in the reclamation period. The following situations show how retention is affected:

- The reclaimable space on VolumeA is less than 60%. The retention date of VolumeA is extended by 365 days.

- The reclaimable space on VolumeA is greater than 60%, and the remaining retention time of the data is more than 365 days. VolumeA is reclaimed, and the retention date of the target volume is set based on the remaining retention period of the data plus 30 days for the reclamation period.
- The reclaimable space on VolumeA is greater than 60%, and the retention time of the data is less than 365 days. VolumeA is reclaimed, and its retention date is set to 365 days, the RETENTIONEXTENSION value, plus 30 days for the reclamation period.

Continuous data protection with the SnapLock feature

If data is stored on a volume with the SnapLock feature enabled, and the data is moved or copied to a non-SnapLock volume, the data loses the unique hardware protection that is provided by NetApp WORM volumes.

The IBM Spectrum Protect™ server allows this type of movement. However, if data is moved from a WORM FILE volume to another type of media, the data might no longer be protected from inadvertent or malicious deletion. If this data is on WORM volumes to meet data retention and protection requirements for legal purposes and is moved to other media, the data might no longer meet those requirements. You must configure your storage pools so that this type of data is kept in storage pools that consist of SnapLock WORM volumes during the entire data retention period.

Setting up SnapLock volumes as IBM Spectrum Protect WORM FILE volumes

To meet strict requirements for archived data, enable the NetApp SnapLock feature.

About this task

When you define or update configurations that involve SnapLock storage pools, ensure that the RECLAMATIONTYPE=SNAPLOCK option is specified for the storage pools that are selected for the NEXTSTGPOOL, RECLAIMSTGPOOL, and COPYSTGPOOLS parameters.

When you configure the storage pools in this way, you help to ensure that your data is properly protected. If you define a next, reclaim, copy storage pool, or active-data pool without selecting the RECLAMATIONTYPE=SNAPLOCK option, the storage pool is not protected. The command succeeds, but a warning message is issued.

Procedure

To set up a SnapLock volume for use as an IBM Spectrum Protect™ WORM FILE volume, complete the following steps:

1. Install and set up SnapLock on the NetApp file server. Ensure that you configure the minimum, maximum, and default retention periods. For instructions, see the NetApp documentation.
2. Install and configure an IBM Spectrum Protect server.
3. Enable archive data retention protection by issuing the SET ARCHIVERETENTIONPROTECTION command:


```
set archiveretentionprotection on
```
4. Set up policy by using the DEFINE COPYGROUP command. Select RETVER and RETMIN values in the archive copy group that meet your requirements for protecting this data in WORM storage. If the RETVER or RETMIN values are not specified, the default management classes values are used.
5. Set up storage by using the DEFINE DEVCLASS command.
 - Use the FILE device class.
 - Specify the DIRECTORY parameter to point to the directory or directories on the SnapLock volumes.
6. Define a storage pool by using the device class that is defined in step 5 by issuing the DEFINE STGPOOL command and specifying the RECLAMATIONTYPE=SNAPLOCK parameter.
7. Update the copy group to point to the storage pool by issuing the UPDATE COPYGROUP command.
8. Use the IBM Spectrum Protect API to archive your objects into the SnapLock storage pool. This feature is not available on standard IBM Spectrum Protect backup-archive clients.

Repairing and recovering data

You can repair damaged data extents in directory-container storage pools and recover lost data after a disaster.

Data extents are part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates. If you have damaged files or directories in your directory-container storage pool, you can repair deduplicated data extents from either the target replication server, the source replication server, or from container-copy storage pool tape volumes.

- Repairing storage pools from a target replication server
If files, directories, or storage pools on a source replication server are damaged, you can repair deduplicated data extents in a directory-container storage pool on the source replication server from a target replication server.
- Repairing storage pools from container-copy storage pool volumes
If files, directories, or storage pools on a source server are damaged, you can repair data extents in a directory-container storage pool on the source server by retrieving the deduplicated data extents from onsite or offsite container-copy storage pool tape volumes.
- Repairing storage pools in an environment with both a replication server and container-copy storage pool volumes
If files, directories, or storage pools on a source server are damaged, you can repair data extents in a directory-container storage pool on the source replication server by retrieving the deduplicated data extents from either the target replication server or from container-copy storage pool tape volumes.
- Repairing storage pools on a target replication server
If files, directories, or storage pools on a target replication server are damaged, you can repair data extents in a directory-container storage pool on the target replication server by retrieving the deduplicated data extents from the source replication server.
- Repairing storage pools after a disaster
You can repair directory-container storage pools and recover their lost data after a disaster.
- Replacing a damaged container-copy storage pool tape volume
If a tape volume that is storing a copy of deduplicated data extents in a container-copy storage pool becomes damaged, you can replace the volume.

Related concepts:

Strategies for disaster protection

Related tasks:

Data protection solutions

Recovering from data loss or system outages

Repairing storage pools from a target replication server

If files, directories, or storage pools on a source replication server are damaged, you can repair deduplicated data extents in a directory-container storage pool on the source replication server from a target replication server.

Before you begin

Complete the following steps:

1. Evaluate your storage environment to determine whether outages, network issues, or hardware failures are causing damage to data or causing the data to appear damaged. If issues in your environment are causing damage to data, identify and resolve the issues.
2. Ensure that enough space is available in the directory-container storage pool for the recovered data. The `PREVIEW=YES` parameter in the `REPAIR STGPOOL` command specifies how much data will be repaired. If not enough space is available, use the `DEFINE STGPOOLDIRECTORY` command to provision space.
3. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - Issue the administrative command, `BACKUP DB`.
4. Review the latest information about repairing and recovering data in technote 2013682.
5. To plan the next steps, review the following restrictions about using the `AUDIT CONTAINER` command.

Attention:

- If you issue the `AUDIT CONTAINER` command with the `ACTION=MARKDAMAGED` setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the `REPAIR STGPOOL` command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:
 - a. Consider running the `AUDIT CONTAINER` command with the `ACTION=SCANALL` setting first. The `ACTION=SCANALL` setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the `REPAIR STGPOOL` command.

- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of damage:

- Minor damage that is caused by accidental deletion of files or directories, overwritten files, accidental changes in file permissions, or disk errors caused by hardware issues.
- Moderate damage that is caused by disk errors or disk mount errors. This type of damage results in the loss of one or more directories, but not a loss of the entire storage pool.

Damaged deduplicated extents are repaired with extents that were protected to the target replication server.

Restriction: You can issue the REPAIR STGPOOL command for a specified storage pool only if you already copied the data to another storage pool on a target replication server by using the PROTECT STGPOOL command.

When you repair a directory-container storage pool from a replication server, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The target replication server is unavailable.
- The target storage pool is damaged.
- A network outage occurs.

Procedure

1. If you suspect minor damage, issue the AUDIT CONTAINER command for the container storage pool at the directory level to identify inconsistencies between the database and the directory-container storage pool. By identifying the damaged data extents in the directory-container storage pool, you can determine which data extents to repair. To conserve time and resources, audit only containers that you suspect are damaged. If you suspect that your directory-container storage pool has more serious damage, issue the AUDIT CONTAINER command at the storage pool level.

For example, to audit a directory, n:\pooldir, in a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1 stgpooledirectory=n:\pooldir
```

To audit a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1
```

The audit process might run for several hours.

2. To repair a directory-container storage pool, issue the REPAIR STGPOOL command and specify the SRCLOCATION=REPLSERVER parameter. For example, to repair a storage pool that is named STGPOOL1 from a replication server, issue the following command:

```
repair stgpool stgpool1 srclocation=replserver
```

When you issue the REPAIR STGPOOL command, the damaged extents are deleted from the volume immediately after they are repaired. The damaged extents are not retained according to the value specified by the REUSEDELAY parameter.

3. Identify any additional damaged extents by issuing the QUERY DAMAGED command.
4. If damage is detected and deduplicated extents cannot be repaired from the replication server, it is still possible that they will be repaired. In some cases, the client node resends data during a backup operation and the damaged extents are repaired. Wait two backup cycles to allow client backup operations to occur. After two backup cycles, complete the following steps:
 - a. To confirm that the damage is repaired, reissue the QUERY DAMAGED command.
 - b. If an entire storage pool directory is damaged, create a new replacement storage pool directory using the DEFINE STGPOOLDIRECTORY command.
 - c. To remove objects that refer to damaged data, issue the AUDIT CONTAINER command and specify the ACTION=REMOVEDAMAGED parameter.
For example, to audit a directory-container storage pool that is named STGPOOL1 and remove damaged objects, issue the following command:

```
audit container stgpool=stgpool1 action=removedamaged
```

- d. Optionally, issue the DELETE STGPOOLDIRECTORY command to delete the empty storage pool directory that you replaced with a new directory in step 4.b.

What to do next

If you continue to detect damaged data over time, issue the AUDIT CONTAINER command for the directory-container storage pool to determine whether there is more widespread damage. For example, to audit a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1
```

Related reference:

AUDIT CONTAINER (Verify the consistency of database information for a directory-container)
DEFINE SCHEDULE (Define a schedule for an administrative command)
QUERY DAMAGED (Query damaged storage pool data)
PROTECT STGPOOL (Protect storage pool data)
REPAIR STGPOOL (Repair a directory-container storage pool)
DEFINE STGPOOLDIRECTORY (Define a storage pool directory)
DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)

Repairing storage pools from container-copy storage pool volumes

If files, directories, or storage pools on a source server are damaged, you can repair data extents in a directory-container storage pool on the source server by retrieving the deduplicated data extents from onsite or offsite container-copy storage pool tape volumes.

Before you begin

Complete the following steps:

1. Evaluate your storage environment to determine whether outages, network issues, or hardware failures are causing damage to data or causing the data to appear damaged. If issues in your environment are causing damage to data, identify and resolve the issues.
2. Ensure that enough space is available in the directory-container storage pool for the recovered data. The PREVIEW=YES parameter in the REPAIR STGPOOL command specifies how much data will be repaired. If the space is insufficient, use the DEFINE STGPOOLDIRECTORY command to provision space.
3. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - o On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - o Issue the administrative command, BACKUP DB.
4. Review the latest information about repairing and recovering data in technote 2013682.
5. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.

Attention:

- o If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:
 - a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.
- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of damage:

- Minor damage that is caused by accidental deletion of files or directories, overwritten files, accidental changes in file permissions, or disk errors caused by hardware issues.
- Moderate damage that is caused by disk errors or disk mount errors. This type of damage results in the loss of one or more directories, but not a loss of the entire storage pool.

Damaged deduplicated extents are repaired with extents that were protected to container-copy storage pools.

Restriction: You can issue the REPAIR STGPOOL command for a specified storage pool only if you already copied the data to container-copy storage pools by using the PROTECT STGPOOL command.

When you repair a directory-container storage pool from container-copy pools, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The container-copy storage pool is unavailable.
- The container-copy storage pool is damaged.
- The container-copy storage pool volumes are unavailable or damaged.

Procedure

1. If you suspect minor damage, issue the AUDIT CONTAINER command for the container storage pool at the directory level to identify inconsistencies between the database and the directory-container storage pool. By identifying the damaged data extents in the directory-container storage pool, you can determine which data extents to repair. To conserve time and resources, audit only containers that you suspect are damaged. If you suspect that your container storage pool has more serious damage, issue the AUDIT CONTAINER command at the storage pool level. For example, to audit a directory, `n:\pooldir`, in a storage pool that is named `STGPOOL1`, issue the following command:

```
audit container stgpool=stgpool1 stgpooledirectory=n:\pooldir
```

To audit a storage pool that is named `STGPOOL1`, issue the following command:

```
audit container stgpool=stgpool1
```

The audit process might run for several hours.

During the repair operation, the server prompts you for the volumes that it requires. In step 3, you will bring the volumes onsite and check them into the library. The required volumes must be brought onsite and checked into the library.

2. To preview the repair operation and generate the list of tape volumes that are needed for the repair operation, issue the REPAIR STGPOOL command and specify the `SRCLLOCATION=LOCAL` and `PREVIEW=YES` parameters. For example, to preview the repair operation for a storage pool that is named `STGPOOL1` from container-copy storage pools, issue the following command:

```
repair stgpool stgpool1 srclocation=local preview=yes
```

The preview process might take some time to finish.

3. If some of the required volumes are offsite, complete the following steps:
 - a. Use the list from the preview operation to determine which volumes need to be brought onsite.
 - b. When the volumes are back onsite, check them into the library by issuing the `CHECKIN LIBVOLUME` command and specifying the `STATUS=PRIVATE` parameter.
 - c. Update the status of the volumes by issuing the `UPDATE STGPOOL` command and specifying the `ACCESS=READWRITE` parameter.

For detailed instructions about the disaster recovery manager (DRM) function, see *Using disaster recovery manager for tape environments (V7.1.1)*.

4. Based on the information that you obtained during the preview operation, ensure that the storage pool contains enough space for the recovered data. If there is not enough space, use the `DEFINE STGPOOLDIRECTORY` command to provision space.
5. To repair the directory-container storage pool, issue the REPAIR STGPOOL command and specify the `SRCLLOCATION=LOCAL` parameter.

For example, to repair a storage pool that is named `STGPOOL1` from a container-copy storage pool, issue the following command:

```
repair stgpool stgpool1 srclocation=local
```

When you issue the REPAIR STGPOOL command, the damaged extents are deleted from the volume immediately after they are repaired. The damaged extents are not retained according to the value specified by the `REUSEDELAY` parameter.

6. Identify any additional damaged extents by issuing the QUERY DAMAGED command.
7. If damage is detected and deduplicated extents cannot be repaired from the container-copy storage pools, it is still possible that they will be repaired. In some cases, the client node resends data during a backup operation and the damaged extents are repaired. Wait two backup cycles to allow client backup operations to occur. After two backup cycles, complete the following steps:
 - a. To confirm that the damage is repaired, reissue the QUERY DAMAGED command.
 - b. If an entire storage pool directory is damaged, create a new replacement storage pool directory using the DEFINE STGPOOLDIRECTORY command.
 - c. To remove objects that refer to damaged data, issue the AUDIT CONTAINER command and specify the ACTION=REMOVEDAMAGED parameter.
For example, to audit a directory-container storage pool that is named STGPOOL1 and remove damaged objects, issue the following command:


```
audit container stgpool=stgpool1 action=removedamaged
```
 - d. Optionally, issue the DELETE STGPOOLDIRECTORY command to delete the empty storage pool directory that you replaced with a new directory in step 7.b.
8. If you repaired an entire storage pool directory, delete the original directory, which is empty and was replaced by a new directory. Delete the original directory by issuing the DELETE STGPOOLDIRECTORY command.

What to do next

If you continue to detect damaged data over time, issue the AUDIT CONTAINER command for the directory-container storage pool to determine whether there is more widespread damage. For example, to audit a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1
```

Related reference:

AUDIT CONTAINER (Verify the consistency of database information for a directory-container)
 DEFINE SCHEDULE (Define a schedule for an administrative command)
 QUERY DAMAGED (Query damaged storage pool data)
 PROTECT STGPOOL (Protect storage pool data)
 REPAIR STGPOOL (Repair a directory-container storage pool)
 DEFINE STGPOOLDIRECTORY (Define a storage pool directory)
 DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)

Repairing storage pools in an environment with both a replication server and container-copy storage pool volumes

If files, directories, or storage pools on a source server are damaged, you can repair data extents in a directory-container storage pool on the source replication server by retrieving the deduplicated data extents from either the target replication server or from container-copy storage pool tape volumes.

Before you begin

Complete the following steps:

1. Evaluate your storage environment to determine whether outages, network issues, or hardware failures are causing damage to data or causing the data to appear damaged. If issues in your environment are causing damage to data, identify and resolve the issues.
2. Ensure that enough space is available in the directory-container storage pool for the recovered data. The PREVIEW=YES parameter in the REPAIR STGPOOL command specifies how much data will be repaired. If the space is insufficient, use the DEFINE STGPOOLDIRECTORY command to provision space.
3. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - o On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - o Issue the administrative command, BACKUP DB.
4. Review the latest information about repairing and recovering data in technote 2013682.
5. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.
Attention:
 - o If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database

size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:

- a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
- b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.
- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of damage:

- Minor damage that is caused by accidental deletion of files or directories, overwritten files, accidental changes in file permissions, or disk errors caused by hardware issues.
- Moderate damage that is caused by disk errors or disk mount errors. This type of damage results in the loss of one or more directories, but not a loss of the entire storage pool.

Damaged deduplicated extents are repaired with extents that were protected to the target replication server or to container-copy storage pools on a source server.

Restriction: You can issue the REPAIR STGPOOL command for a specified storage pool only if you already copied the data to another storage pool on a target replication server or to container-copy storage pools by using the PROTECT STGPOOL command. When you repair a directory-container storage pool from a target replication server, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The target replication server is unavailable.
- The target storage pool is damaged.
- A network outage occurs.

When you repair a directory-container storage pool from container-copy pools, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The container-copy storage pool is unavailable.
- The container-copy storage pool is damaged.
- The container-copy storage pool volumes are unavailable or damaged.

Procedure

1. Attempt to repair the storage pool from the target replication server by completing the steps in Repairing storage pools from a target replication server.
2. If the damaged extents cannot be repaired from the target replication server, repair the damaged extents from container-copy storage pools by completing the steps in Repairing storage pools from container-copy storage pool volumes.
3. If you repaired damaged extents from container-copy storage pools, issue the PROTECT STGPOOL command and specify the TYPE=REPLSERVER parameter for the storage pools on the source replication server.

What to do next

If you continue to detect damaged data over time, issue the AUDIT CONTAINER command for the directory-container storage pool to determine whether there is more widespread damage. For example, to audit a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1
```

Related reference:

AUDIT CONTAINER (Verify the consistency of database information for a directory-container)

DEFINE SCHEDULE (Define a schedule for an administrative command)

QUERY DAMAGED (Query damaged storage pool data)

PROTECT STGPOOL (Protect storage pool data)
REPAIR STGPOOL (Repair a directory-container storage pool)
DEFINE STGPOOLDIRECTORY (Define a storage pool directory)
DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)

Repairing storage pools on a target replication server

If files, directories, or storage pools on a target replication server are damaged, you can repair data extents in a directory-container storage pool on the target replication server by retrieving the deduplicated data extents from the source replication server.

Before you begin

Complete the following steps:

1. Evaluate your storage environment to determine whether outages, network issues, or hardware failures are causing damage to data or causing the data to appear damaged. If issues in your environment are causing damage to data, identify and resolve the issues.
2. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - o On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - o Issue the administrative command, BACKUP DB.
3. Review the latest information about repairing and recovering data in technote 2013682.
4. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.

Attention:

- o If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:
 - a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.
- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of damage:

- Minor damage that is caused by accidental deletion of files or directories, overwritten files, accidental changes in file permissions, or disk errors caused by hardware issues.
- Moderate damage that is caused by disk errors or disk mount errors. This type of damage results in the loss of one or more directories, but not a loss of the entire storage pool.

As part of the operation of the PROTECT STGPOOL command, damaged extents in the target storage pool are repaired. To be repaired, extents must already be marked as damaged on the target server. For example, an AUDIT CONTAINER command might identify damage in the target storage pool before the PROTECT STGPOOL command is issued.

Procedure

1. Protect data extents in a directory-container storage pool on a source server by issuing the PROTECT STGPOOL command. For example, to protect a directory-container storage pool that is named POOL1, issue the following command:

```
protect stgpool pool1
```

Wait for the protection process to finish.

2. To identify the damaged data extents in the directory-container storage pool on the target server, issue the AUDIT CONTAINER command.

For example, to audit a storage pool that is named STGPOOL1, issue the following command:

```
audit container stgpool=stgpool1
```

3. Repair damaged extents in the target storage pool by reissuing the PROTECT STGPOOL command on the source server. The damaged extents in the target storage pool are marked as damaged and are repaired.
4. Confirm that there are no additional damaged extents by issuing the QUERY DAMAGED command.

Related reference:

AUDIT CONTAINER (Verify the consistency of database information for a directory-container)

DEFINE SCHEDULE (Define a schedule for an administrative command)

QUERY DAMAGED (Query damaged storage pool data)

PROTECT STGPOOL (Protect storage pool data)

REPAIR STGPOOL (Repair a directory-container storage pool)

DEFINE STGPOOLDIRECTORY (Define a storage pool directory)

DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)

Repairing storage pools after a disaster

You can repair directory-container storage pools and recover their lost data after a disaster.

If a disaster occurs and your primary site is no longer available, you can repair your directory-container storage pools by restoring them on a new target server at your recovery site.

- Repairing storage pools from container-copy storage pool volumes after a disaster
If a disaster occurs on a source server, you can repair deduplicated data extents in a directory-container storage pool from offsite container-copy storage pool tape volumes. The directory-container storage pool is repaired on a target server at a recovery site.
- Repairing storage pools from a target replication server after a disaster
If a disaster occurs on a source replication server, you can repair deduplicated data extents in a directory-container storage pool from a target replication server. The directory-container storage pool is repaired on a target server at a recovery site.
- Repairing storage pools in an environment with both a replication server and container-copy storage pool volumes after a disaster
If a disaster occurs on a source server, you can repair deduplicated data extents in a directory-container storage pool from a replication target server or from offsite container-copy storage pool tape volumes. The directory-container storage pool is repaired on a target server at a recovery site.

Related reference:

Determining whether to use container-copy storage pools for disaster protection

Repairing storage pools from container-copy storage pool volumes after a disaster

If a disaster occurs on a source server, you can repair deduplicated data extents in a directory-container storage pool from offsite container-copy storage pool tape volumes. The directory-container storage pool is repaired on a target server at a recovery site.

Before you begin

Complete the following steps:

1. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - Issue the administrative command, BACKUP DB.
2. Review the latest information about repairing and recovering data in technote 2013682.
3. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.

Attention:

 - If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or

days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:

- a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.
- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
- a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of major damage:

- Complete loss of all container storage pools on the source server
- Complete loss of the primary site

The following assumptions are made for this disaster recovery scenario:

- You were using the PROTECT STGPOOL command to back up data to offsite container-copy storage pools from a source server. You retrieved the offsite tape volumes and have them at your recovery site.
- You were not using the PROTECT STGPOOL command to back up data to a target replication server.
- You used the IBM Spectrum Protect Blueprints to configure the IBM Spectrum Protect source server, and you also used the Blueprint configuration scripts to restore the environment by setting up a new target server at a recovery site. The scripts copied backup versions of the IBM Spectrum Protect database, the server options file (dsmserv.opt), the volume history file (volhist.out), and device configuration file (devconfig.out) to their original locations on the recovery server. After the scripts run, you see the newly created, empty directories on the recovery server.

When you attempt to repair a directory-container storage pool from container-copy pools, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The container-copy storage pool is unavailable.
- The container-copy storage pool is damaged.
- The container-copy storage pool volumes are unavailable or damaged.

Procedure

1. Mark all data extents in the container storage pool as damaged by issuing the AUDIT CONTAINER command for the container storage pool at the storage pool level, and specifying the ACTION=MARKDAMAGED parameter. For example, to audit a storage pool that is named STGPOOL1 and mark it as damaged, issue the following command:

```
audit container stgpool=stgpool1 action=markdamaged
```

2. If you protected the directory-container storage pool by using both onsite and offsite container-copy storage pools, issue the UPDATE STGPOOL command for the onsite copy of the container-copy storage pools, and specify the ACCESS=UNAVAILABLE parameter.
3. When the offsite container-copy storage pool volumes are back onsite, check them into the library by issuing the CHECKIN LIBVOLUME command and specifying the STATUS=PRIVATE parameter.
4. Update the status of the volumes by issuing the UPDATE STGPOOL command and specifying the ACCESS=READWRITE parameter.
5. Repair the storage pool by issuing the REPAIR STGPOOL command and specifying the SRCLOCATION=LOCAL parameter. For example, to repair a storage pool that is named STGPOOL1 from offsite container-copy storage pools, issue the following command:

```
repair stgpool stgpool1 srclocation=local
```

When you issue the REPAIR STGPOOL command, the damaged extents are deleted from the volume immediately after they are repaired. The damaged extents are not retained according to the value specified by the REUSEDELAY parameter.

6. Confirm that there are no additional damaged extents by issuing the QUERY DAMAGED command.

7. Repeat this procedure to repair all of your storage pools.

Repairing storage pools from a target replication server after a disaster

If a disaster occurs on a source replication server, you can repair deduplicated data extents in a directory-container storage pool from a target replication server. The directory-container storage pool is repaired on a target server at a recovery site.

Before you begin

Complete the following steps:

1. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - o On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - o Issue the administrative command, BACKUP DB.
2. Review the latest information about repairing and recovering data in technote 2013682.
3. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.

Attention:

 - o If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:
 - a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.
 - o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of major damage:

- Complete loss of all container storage pools on the source replication server
- Complete loss of the primary site

The following assumptions are made for this disaster recovery scenario:

- You were using the PROTECT STGPOOL command to back up data from a source replication server to a target replication server. The target replication server is running at your recovery site.
- You were not using the PROTECT STGPOOL command to back up data to offsite container-copy storage pools.
- You used the IBM Spectrum Protect Blueprints to configure the IBM Spectrum Protect source server, and you also used the Blueprint configuration scripts to restore the environment by setting up a new target server at a recovery site. The scripts copied backup versions of the IBM Spectrum Protect database, the server options file (dsmserv.opt), the volume history file (volhist.out), and device configuration file (devconfig.out) to their original locations on the recovery server. After the scripts run, you see the newly created, empty directories on the recovery server.

When you attempt to repair a directory-container storage pool from a target replication server, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The target replication server is unavailable.
- The target storage pool is damaged.
- A network outage occurs.

Procedure

1. Mark all data extents in the container storage pool as damaged by issuing the AUDIT CONTAINER command for the container storage pool at the storage pool level, and specifying the ACTION=MARKDAMAGED parameter. For example, to audit a storage pool that is named STGPOOL1 and mark it as damaged, issue the following command:

```
audit container stgpool=stgpool1 action=markdamaged
```

2. Repair the storage pool by issuing the REPAIR STGPOOL command and specifying the SRCLOCATION=REPLSERVER parameter. For example, to repair a storage pool that is named STGPOOL1 from a target replication server, issue the following command:

```
repair stgpool stgpool1 srclocation=replserver
```

When you issue the REPAIR STGPOOL command, the damaged extents are deleted from the volume immediately after they are repaired. The damaged extents are not retained according to the value specified by the REUSEDELAY parameter.

3. If you did not use Blueprint configuration scripts to set up your target replication server, the file structure on the target replication server might not match the information that is stored in the database. Optionally, remove the storage pool directories that do not exist on the target replication server by issuing the DELETE STGPOOLDIRECTORY command.
4. Confirm that there are no additional damaged extents by issuing the QUERY DAMAGED command.
5. If damage is detected and deduplicated extents cannot be repaired from the replication server, it is still possible that they will be repaired. In some cases, the client node resends data during a backup operation and the damaged extents are repaired. Wait two backup cycles to allow client backup operations to occur. After two backup cycles, complete the following steps:
 - a. To confirm that the damage is repaired, reissue the QUERY DAMAGED command.
 - b. To remove objects that refer to damaged data, issue the AUDIT CONTAINER command and specify the ACTION=REMOVEDAMAGED parameter. For example, to audit a directory-container storage pool that is named STGPOOL1 and remove damaged objects, issue the following command:

```
audit container stgpool=stgpool1 action=removedamaged
```

6. Repeat this procedure to repair all of your storage pools.

Related reference:

QUERY DAMAGED (Query damaged storage pool data)

Repairing storage pools in an environment with both a replication server and container-copy storage pool volumes after a disaster

If a disaster occurs on a source server, you can repair deduplicated data extents in a directory-container storage pool from a replication target server or from offsite container-copy storage pool tape volumes. The directory-container storage pool is repaired on a target server at a recovery site.

Before you begin

Complete the following tasks:

1. Back up the IBM Spectrum Protect™ server database by using one of the following methods:
 - o On the Operations Center Overviews page, click Servers, select a server, and click Back Up.
 - o Issue the administrative command, BACKUP DB.
2. Review the latest information about repairing and recovering data in technote 2013682.
3. To plan the next steps, review the following restrictions about using the AUDIT CONTAINER command.

Attention:

 - o If you issue the AUDIT CONTAINER command with the ACTION=MARKDAMAGED setting for an entire storage pool, referenced data is unavailable for restore operations until the storage pool is repaired. Depending on the database size, network bandwidth, media speed, and other factors, the REPAIR STGPOOL command might run for hours or days. For this reason, if some of the data in the storage pool is available, or the status of data in the storage pool is unknown, follow these guidelines:
 - a. Consider running the AUDIT CONTAINER command with the ACTION=SCANALL setting first. The ACTION=SCANALL setting identifies database records that refer to data extents with inconsistencies. Only those data extents are marked as damaged in the database.
 - b. After the extents are marked as damaged, you can run the REPAIR STGPOOL command.

- o If you plan to run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting, follow these guidelines:
 - a. Considering running the QUERY DAMAGED command first to determine the scope of damaged data extents in the storage pool.
 - b. After that, you can run the REPAIR STGPOOL command to repair damaged extents in the storage pool.
 - c. Finally, you can run the AUDIT CONTAINER command with the ACTION=REMOVEDAMAGED setting to remove any damaged data extents that remain in the storage pool.

About this task

Use the procedure to repair the following types of major damage:

- Complete loss of all container storage pools on the source server
- Complete loss of the primary site

The following assumptions are made for this disaster recovery scenario:

- You were using the PROTECT STGPOOL command to back up data from a source replication server to a target replication server. The target replication server is running at your recovery site.
- You were using the PROTECT STGPOOL command to back up data to offsite container-copy storage pools.
- You used the IBM Spectrum Protect Blueprints to configure the IBM Spectrum Protect source server, and you also used the Blueprint configuration scripts to restore the environment by setting up a new target server at a recovery site. The scripts copied backup versions of the IBM Spectrum Protect database, the server options file (dsmserv.opt), the volume history file (volhist.out), and device configuration file (devconfig.out) to their original locations on the recovery server. After the scripts run, you see the newly created, empty directories on the recovery server.

When you attempt to repair a directory-container storage pool from a target replication server, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The target replication server is unavailable.
- The target storage pool is damaged.
- A network outage occurs.

When you repair a directory-container storage pool from container-copy pools, the REPAIR STGPOOL command fails if any of the following conditions occur:

- The container-copy storage pool is unavailable.
- The container-copy storage pool is damaged.
- The container-copy storage pool volumes are unavailable or damaged.

Procedure

1. Mark all data extents in the container storage pool as damaged by issuing the AUDIT CONTAINER command for the container storage pool at the storage pool level, and specifying the ACTION=MARKDAMAGED parameter. For example, to audit a storage pool that is named STGPOOL1 and mark it as damaged, issue the following command:

```
audit container stgpool=stgpool1 action=markdamaged
```

2. If you protected the directory-container storage pool by using both onsite and offsite container-copy storage pools, issue the UPDATE STGPOOL command for the onsite copy of the container-copy storage pools, and specify the ACCESS=UNAVAILABLE parameter.
3. When the offsite container-copy storage pool volumes are back onsite, check them into the library by issuing the CHECKIN LIBVOLUME command and specifying the STATUS=PRIVATE parameter. By moving the tape volumes onsite now, you are prepared to repair damaged extents from the container-copy tape volumes if the damaged extents cannot be repaired from the target replication server.
4. Update the status of the volumes by issuing the UPDATE STGPOOL command and specifying the ACCESS=READWRITE parameter.
5. Repair the storage pool by issuing the REPAIR STGPOOL command and specifying the SRCLOCATION=REPLSERVER parameter. For example, to repair a storage pool that is named STGPOOL1 from a target replication server, issue the following command:

```
repair stgpool stgpool1 srclocation=replserver
```

When you issue the REPAIR STGPOOL command, the damaged extents are deleted from the volume immediately after they are repaired. The damaged extents are not retained according to the value specified by the REUSEDELAY parameter.

6. If you did not use Blueprint configuration scripts to set up your target replication server, the file structure on the target replication server might not match the information that is stored in the database. Optionally, remove the storage pool directories that do not exist on the target replication server. Issue the DELETE STGPOOLDIRECTORY command to delete directories that are not on the target replication server.
7. Confirm that there are no additional damaged extents by issuing the QUERY DAMAGED command.
8. If the damaged extents cannot be repaired from the target replication server, you can repair the damaged extents from offsite container-copy storage pools. For instructions, see [Repairing storage pools from container-copy storage pool volumes after a disaster](#).
9. Confirm that there are no additional damaged extents by reissuing the QUERY DAMAGED command.
10. If damage is detected and deduplicated extents cannot be repaired from the replication server, it is still possible that they will be repaired. In some cases, the client node resends data during a backup operation and the damaged extents are repaired. Wait two backup cycles to allow client backups to occur. After two backup cycles, complete the following steps:
 - a. To confirm that the damage is repaired, reissue the QUERY DAMAGED command.
 - b. To remove objects that refer to damaged data, issue the AUDIT CONTAINER command and specify the ACTION=REMOVEDAMAGED parameter.
For example, to audit a directory-container storage pool that is named STGPOOL1 and remove damaged objects, issue the following command:

```
audit container stgpool=stgpool1 action=removedamaged
```

11. Repeat this procedure to repair all of your storage pools.

Replacing a damaged container-copy storage pool tape volume

If a tape volume that is storing a copy of deduplicated data extents in a container-copy storage pool becomes damaged, you can replace the volume.

Procedure

1. Delete the damaged tape volume by issuing the DELETE VOLUME command and specifying the DISCARDATA=YES parameter.

For example, to delete a volume that is named VOLUME1, issue the following command:

```
delete volume volume1 discarddata=yes
```

2. Protect data extents in the directory-container storage pool by copying the data to existing volumes in the container-copy storage pool. Issue the PROTECT STGPOOL command from the source server.

For example, to protect a directory-container storage pool that is named POOL1, issue the following command:

```
protect stgpool pool1 type=local
```

Related reference:

PROTECT STGPOOL (Protect storage pool data)

DELETE VOLUME (Delete a storage pool volume)

Server commands, options, and utilities

Use commands to administer and configure the server, options to customize the server, and utilities to perform special tasks when the server is not running.

- **Managing the server from the command line**
IBM Spectrum Protect™ provides several different command-line interfaces for managing IBM Spectrum Protect servers.
- **Administrative commands**
Administrative commands are available to manage and configure the server.
- **Server options**
At installation, IBM Spectrum Protect provides a server options file that contains a set of default options to start the server.
- **Server utilities**
Use server utilities to perform special tasks on the server while the server is not running.
- **Return codes for use in IBM Spectrum Protect scripts**
You can write IBM Spectrum Protect scripts that use return codes to determine how script processing proceeds. The return

codes can be one of three severities: OK, WARNING, ERROR.

- Device utilities
You can use device utilities for tasks that are related to configuring storage devices for the IBM Spectrum Protect server.
- Server scripts and macros for automation
You can automate common administrative tasks by creating IBM Spectrum Protect server scripts or administrative client macros. Server scripts are stored in the server database and can be scheduled to run with an administrative schedule command. Administrative client macros are stored as files on the administrative client.

Managing the server from the command line

IBM Spectrum Protect™ provides several different command-line interfaces for managing IBM Spectrum Protect servers.

About this task

The following command-line interfaces are available:

Administrative command-line client

The administrative command-line client is a program that runs on a file server, workstation, or mainframe. It is installed as part of the IBM Spectrum Protect server installation process. The administrative client can be accessed remotely.

From the administrative client, you can issue any server commands.

Server console

The server console is a command-line window on the system where the server is installed. Therefore, to use the server console, you must be at the physical location of the server system.

Compared to the administrative client, the capabilities of the server console are limited. From the server console, you cannot issue certain commands, and you cannot route commands to other servers. Also, you cannot specify that certain commands process before other commands can be issued. However, this limitation can be useful if, for example, you want to run two commands in quick succession.

Operations Center command line

From the Operations Center, you can access the IBM Spectrum Protect command line. You might want to use this command line to issue server commands to complete certain IBM Spectrum Protect tasks that are not supported in the Operations Center.

Server scripts provide for automation of common administrative tasks. A macro is a file that contains one or more IBM Spectrum Protect administrative commands. When you issue the MACRO command, the server processes all commands in the macro file in order, including commands that are contained in any nested macros.

- Issuing commands from the administrative client
The administrative command-line client is a program that runs on a file server, workstation, or mainframe.
- Issuing commands from the Operations Center
From the Operations Center command-line interface, you can issue commands to manage IBM Spectrum Protect servers that are configured as hub or spoke servers.
- Issuing commands from the server console
IBM Spectrum Protect provides a user ID named SERVER_CONSOLE that allows you to issue commands and administer the server from the server console after IBM Spectrum Protect is installed. At installation, SERVER_CONSOLE is automatically registered as an administrator and is given system authority.
- Entering administrative commands
Commands consist of command names and usually parameters and variables. Syntax diagrams depict the rules to follow when entering commands.
- Controlling command processing
You can run some IBM Spectrum Protect commands sequentially or concurrently with other commands. You can also route commands from one server to other servers for processing.
- Performing tasks concurrently on multiple servers
Command routing allows you to route commands to one or more servers for processing and then collect the output from these servers.
- Privilege classes for commands
The authority granted to an administrator through the privilege class determines which administrative commands that the administrator can issue.

Related concepts:

Issuing commands from the administrative client

The administrative command-line client is a program that runs on a file server, workstation, or mainframe.

About this task

Ensure that your administrative client and your server are running in compatible languages. See LANGUAGE for language and locale options. If your client and server are using different languages, the messages that IBM Spectrum Protect™ generates might not be understandable.

Tip: Text strings that are sent from the client to the server do not depend on the server language setting. The text is displayed properly if the administrative client runs in the same locale when sending the string and when receiving the string.

For example, assume that you update a node contact field with a value that contains national characters (`update node myNode contact=NLcontact_info`), and later query the node (`query node myNode format=detailed`). If the client is running in the same locale when you update as when you query, the `NLcontact_info` displays properly. If you update the node contact field when the client is running in one locale, and query the node when the client is running in a different locale, the `NLcontact_info` might not display properly.

- Starting and stopping the administrative client
Use the DSMADMC command to start an administrative client session.
- Monitoring server activities from the administrative client
To monitor IBM Spectrum Protect activities, such as server migration and client logons, run the administrative client in console mode. You cannot enter any administrative commands in console mode.
- Monitoring removable-media mounts from the administrative client
To monitor the mounting and dismounting of removable media, run the administrative client in mount mode. When the client is running in mount mode, you cannot enter any administrative commands.
- Processing individual commands from the administrative client
Use batch mode to enter a single administrative command. Your administrative client session automatically ends when the command is processed.
- Processing a series of commands from the administrative client
Use the interactive mode to process a series of administrative commands.
- Formatting output from commands
IBM Spectrum Protect formats the output processed from commands according to your screen or window width.
- Saving command output to a specified location
The most common use for redirecting output is to save the output from query commands to a specified file or program. You can then browse the contents of the file or in some cases, print the contents.
- Administrative client options
In all administrative client modes, you can use options to modify administrative client session responses.

Starting and stopping the administrative client

Use the DSMADMC command to start an administrative client session.

About this task

The IBM Spectrum Protect™ server must be running before an administrative client can connect.

Procedure

- To start an administrative client session in command-line mode, enter this command on your workstation:

```
dsmadmc -id=admin -password=adminpwd -dataonly=yes
```

By entering the DSMADMC command with the `-ID` and `-PASSWORD` options as shown, you are not prompted for a user ID and password.

- To stop an administrative command-line client session, enter the following command:

quit

- To interrupt a DSMADMC command before the IBM Spectrum Protect server finishes processing it, use the UNIX `kill -9` command from an available command line. Do not press `Ctrl+C` because, while it ends the session, it can lead to unexpected results.

Monitoring server activities from the administrative client

To monitor IBM Spectrum Protect™ activities, such as server migration and client logons, run the administrative client in console mode. You cannot enter any administrative commands in console mode.

Procedure

- To start an administrative client session in console mode, enter the following command:

```
dsmadmc -consolemode
```

You are prompted for a password if authentication is turned on for the server. If you do not want to be prompted for your user ID and password, enter the DSMADMC command with the `-ID` and `-PASSWORD` options.

- To end an administrative client session in console mode, use a keyboard break sequence.

Operating system	Break sequence
UNIX and Linux clients	Ctrl+C
Windows clients	Ctrl+C or Ctrl+Break

Monitoring removable-media mounts from the administrative client

To monitor the mounting and dismounting of removable media, run the administrative client in mount mode. When the client is running in mount mode, you cannot enter any administrative commands.

Procedure

- To start an administrative client session in mount mode, enter the following command:

```
dsmadmc -mountmode
```

You are prompted for a password if authentication is turned on for the server. If you do not want to be prompted for your user ID and password, enter the DSMADMC command with the `-ID` and `-PASSWORD` options.

- To end an administrative client session in mount mode, use a keyboard break sequence.

Operating system	Break sequence
UNIX and Linux clients	Ctrl+C
Windows clients	Ctrl+C or Ctrl+Break

Processing individual commands from the administrative client

Use batch mode to enter a single administrative command. Your administrative client session automatically ends when the command is processed.

Procedure

To start an administrative client session in batch mode, use the following command: `dsmadmc server_command`

If you do not want to be prompted for your user ID and password, you can enter the DSMADMC command with the `-ID` and `-PASSWORD` options.

In batch mode, you must enter the complete command on one line. If a command does not fit on one line, enter the command by using a macro or a script. If you specify a parameter with a string of text using batch mode, enclose the text in single quotation

marks (') in the macro. Do not use double quotation marks for commands in batch mode, because your operating system might not parse the quotation marks correctly.

Windows You can bypass this batch mode double quotation mark restriction for Windows clients by using the back slash (\) escape character. For example, on the OBJECTS parameter of the DEFINE CLIENTACTION command, you could enter the string with the \ character preceding the double quotation marks in the command.

```
dsmadmc -id=admin -password=admin define clientaction test_node domain=test_dom  
action=restore objects='\"C:\program files\test\*\"'
```

Processing a series of commands from the administrative client

Use the interactive mode to process a series of administrative commands.

About this task

To start an administrative client session in interactive mode, a server session must be available. To ensure the availability of server sessions for both administrative and client node sessions, the interactive mode of the administrative client is disconnected if one or more of the following conditions is true:

- The server was stopped by using the HALT command.
- Commands were not issued from the administrative client session for the length of time that is specified with the IDLETIMEOUT server option.
- The administrative client session was canceled with the CANCEL SESSION command.

Procedure

To start an administrative session in interactive mode, use the following command: `dsmadmc`

You can use continuation characters when you use interactive mode. For more information, see `t_cmdline_longcmd.dita#t_cmdline_longcmd`.

You can automatically restart your administrative client session by entering another command each time the `tsm: servername >` prompt appears.

Do not enter a server command with the DSMADMC command. Doing so starts the administrative client in batch, not interactive, mode. For example, do not enter:

```
dsmadmc server_command
```

Formatting output from commands

IBM Spectrum Protect™ formats the output processed from commands according to your screen or window width.

Procedure

- If the width of your screen or window is not wide enough to display the output horizontally, IBM Spectrum Protect arranges and displays the information vertically.
- You can format the output of QUERY commands using the DISPLAYMODE and OUTFILE administrative client options.

Saving command output to a specified location

The most common use for redirecting output is to save the output from query commands to a specified file or program. You can then browse the contents of the file or in some cases, print the contents.

About this task

On some operating systems, you can redirect output of a command by using special characters such as >, >>, and |. Redirection characters direct the output of a command to a file or program that you specify instead of to your screen. You can save the output from a command by entering redirection characters at the end of the command. To redirect output, leave a blank between the redirection character and the file or program name. See the following examples.

When redirecting output, follow the naming conventions of the operating system where you are running the administrative client.

Procedure

The examples in the following table show how to redirect command output.

Task	Procedure
Redirect the output of a QUERY DOMAIN command to a new file in batch or interactive mode	Use a single greater-than sign (>) to redirect the output to a new file or write over an existing file: <code>dsmadmc -id=sullivan -pa=secretpwd query domain acctg > dominfo.acc</code>
Append the output of a QUERY DOMAIN command to the end of an existing file in batch or interactive mode	Use two consecutive greater-than signs (>>) to append the output to the end of an existing file: <code>dsmadmc -id=sullivan -pa=secretpwd query domain acctg >> dominfo.acc</code>
Redirect all output from an administrative client session in console mode to a program called filter.exe	Use the vertical bar () to direct all output for a session to a program: <code>dsmadmc -console -id=sullivan -password=secretpwd filter.exe</code> The program can be set up to monitor the output for individual messages as they occur and take appropriate action, such as sending mail to another user.
In console mode, redirect all output to a file	Specify the -OUTFILE option with a destination file name. For example, the following command redirects all output to the save.out file: <code>dsmadmc -id=sullivan -password=secretpwd -consolemode -outfile=save.out</code>

Administrative client options

In all administrative client modes, you can use options to modify administrative client session responses.

Syntax

```
          .------.
          v               |
>>-DSMADMC-----+-----+-----+-----+----->>
                '-admin_client_option-'   '-server_command-'
```

Example of using administrative client options

You can enter the DSMADMC command with your user ID and password by using the -ID and -PASSWORD options so that you are not prompted for that information. To have IBM Spectrum Protect™ redirect all output to a file, specify the -OUTFILE option with a destination file name. For example, to issue the QUERY NODE command in batch mode with the output redirected to the SAVE.OUT file, enter:

```
dsmadmc -id=sullivan -password=secret -outfile=save.out query node
```

Options

Administrative client options can be specified with the DSMADMC command and are valid from an administrative client session only. You can type an option in uppercase letters, lowercase letters, or any combination. Uppercase letters denote the shortest acceptable abbreviation. If an option appears entirely in uppercase letters, you cannot abbreviate it.

-ALWAYSPrompt

Specifies that a command prompt is displayed if the input is from the keyboard or if it is redirected (for example, from a file). If this option is not specified and the input is redirected, the command prompt is not written.

If the input is redirected, only the command output is displayed. If this option is specified, the command prompt and the command output are displayed.

-CHECKAliashalt

Allows the administrative client to recognize an alias for the HALT command as set in the ALIASHALT server option. See ALIASHALT for details.

-COMMA delimited

Specifies that any tabular output from a server query is to be formatted as comma-separated strings rather than in readable format. This option is intended to be used primarily when you redirect the output of an SQL query (SELECT command). The comma-separated value format is a standard data format, which can be processed by many common programs, including spreadsheets, databases, and report generators.

-CONSOLE mode

Specifies that IBM Spectrum Protect runs in console mode. Most server console output is echoed to your screen. The exception are items such as responses to query commands that are issued from the console, trace output, or any system messages that displayed on the console.

-DATAONLY=NO or YES

Specifies whether product version information and output headers display with the output. The default is NO.

NO

Specifies that the product version information and output column headers display.

YES

Suppresses the product version information and output column headers.

-DISPLAY mode=LIST or TABLE

You can force the QUERY output to tabular or list format regardless of the command-line window column width.

If you are using the -DISPLAYMODE option and you want the output to go to a file, do not specify the -OUTFILE option. Use redirection to write to the file.

-ID=userid

Specifies the administrator's user ID.

-ITEM commit

Specifies that IBM Spectrum Protect commits commands inside a script or a macro as each command is processed.

-MOUNT mode

Specifies that IBM Spectrum Protect runs in mount mode. All server removable-media mount messages are echoed to your screen.

-NEWLINE AFTER Prompt

Specifies that a newline character is written after the command prompt and commands that are entered from the keyboard are displayed underneath the prompt. If this option is not specified, commands entered from the keyboard are displayed to the right side of the prompt.

-NO Confirm

Specifies that you do not want IBM Spectrum Protect to request confirmation before processing commands that affect the availability of the server or data that is managed by the server.

-OUT file

Specifies that output from a server query is displayed in one row. If the output in a row exceeds the column width that is defined by the server, the output is displayed on multiple lines in that row. This option is available in batch mode only.

-OUT file=filename

Specifies that output from a server query is redirected to a specified file. In batch mode, output is redirected to a file you specify and the format of the output matches the format of the output on your screen.

In interactive, console, or mount mode sessions, output displays on your screen.

-PASSWORD=password

Specifies the administrator's password.

-Quiet

Specifies that IBM Spectrum Protect does not display standard output messages to your screen. However, when you use this option, certain error messages still appear.

AIX Linux -SERVER address

AIX Linux Specifies the server stanza in the dsm.sys file. The client uses the server stanza to determine the server it connects to. The SERVERADDRESS option is supported by administrative clients that are running on UNIX, Linux, and Macintosh operating systems only.

-TAB delimited

Specifies that any tabular output from a server query is to be formatted as tab-separated strings rather than in readable format. This option is intended to be used primarily when you redirect the output of an SQL query (SELECT command). The tab-separated value format is a standard data format, which can be processed by many common programs, including spreadsheets, databases, and report generators.

-TCP Port

Specifies a TCP/IP port address for an IBM Spectrum Protect server. The TCPPOINT option is only supported by administrative clients that are running on Windows operating systems and is valid on the Windows administrative client command line.

-TCPServeraddress

Specifies a TCP/IP server address for an IBM Spectrum Protect server. The TCPSEVERADDRESS option is only supported by administrative clients that are running on Windows operating systems and is valid on the Windows administrative client command line.

In addition to the options that are listed here, you can also specify any option that is in the client options file. Each option must be preceded with a hyphen and delimited with a space.

Issuing commands from the Operations Center

From the Operations Center command-line interface, you can issue commands to manage IBM Spectrum Protect™ servers that are configured as hub or spoke servers.

Procedure

To open the command-line interface, hover over the globe icon  in the Operations Center menu bar, and click Command Builder.

Issuing commands from the server console

IBM Spectrum Protect™ provides a user ID named SERVER_CONSOLE that allows you to issue commands and administer the server from the server console after IBM Spectrum Protect is installed. At installation, SERVER_CONSOLE is automatically registered as an administrator and is given system authority.

About this task

If you have system privilege, you can revoke or grant new privileges to the SERVER_CONSOLE user ID. You cannot take any of the following actions:

- Register or update the SERVER_CONSOLE user ID
- Lock or unlock the SERVER_CONSOLE user ID
- Rename the SERVER_CONSOLE user ID
- Remove SERVER_CONSOLE user ID
- Route commands from the SERVER_CONSOLE user ID

Not all IBM Spectrum Protect commands are supported by the server console. You cannot specify the WAIT parameter from the server console.

Entering administrative commands

Commands consist of command names and usually parameters and variables. Syntax diagrams depict the rules to follow when entering commands.

About this task

To display command-line help for server commands that have unique names, you can type `help commandName`, where *commandName* is the name of the server command for which you want information. For example, to display help for the REGISTER NODE command, type `help register node`. Command syntax and parameter descriptions are displayed in the output.

You can also type `help` followed by the topic number for the command. Topic numbers are listed in the table of contents for command-line help, for example:

```
3.0 Administrative commands
  3.46 REGISTER
    3.46.1 REGISTER ADMIN (Register an administrator)
    3.46.2 REGISTER LICENSE (Register a new license)
    3.46.3 REGISTER NODE (Register a node)
```

To display help about the REGISTER NODE command, type:

```
help 3.46.3
```

Use topic numbers to display command-line help for subcommands. DEFINE DEVCLASS is an example of a command that has subcommands. For example, you can specify the DEFINE DEVCLASS command for 3590 device classes and for 3592 device classes:

```
3.0 Administrative commands
...
3.13.10 DEFINE DEVCLASS (Define a device class)
    3.13.10.1 DEFINE DEVCLASS (Define a 3590 device class)
    3.13.10.2 DEFINE DEVCLASS (Define a 3592 device class)
    ...
```

To display help for the DEFINE DEVCLASS command for 3590 device classes, type:

```
help 3.13.10.1
```

- Reading syntax diagrams
To read a syntax diagram for entering a command, follow the path of the line. Read from left to right and from top to bottom.
- Using continuation characters to enter long commands
Continuation characters are useful when you want to process a command that is longer than your screen or window width. You can use continuation characters in the interactive mode of the administrative client.
- Naming IBM Spectrum Protect objects
IBM Spectrum Protect restricts the number and type of characters that you can use to name objects.
- Using wildcard characters to specify object names
In some commands, such as the query commands, you can use wildcard characters to create a pattern-matching expression that specifies more than one object. Using wildcard characters makes it easier to tailor a command to your needs.
- Specifying descriptions in keyword parameters
If a description (a string of text) for a parameter begins with a single or double quotation mark, or contains any embedded blanks or equal signs, you must surround the value with either single (') or double (") quotation marks.

Reading syntax diagrams

To read a syntax diagram for entering a command, follow the path of the line. Read from left to right and from top to bottom.

- The >>--- symbol indicates the beginning of a syntax diagram.
- The ---> symbol at the end of a line indicates that the syntax diagram continues onto the next line.
- The >--- symbol at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The --->< symbol indicates the end of a syntax diagram.

Command names

The command name can consist of a single action word, such as HALT, or it can consist of an action word and an object for the action, such as DEFINE DOMAIN. You can enter the command in any column of the input line.

Enter the entire command name or the abbreviation that is specified in the syntax diagram for the command. Uppercase letters denote the shortest acceptable abbreviation. If a command appears entirely in uppercase letters, you cannot abbreviate it. You can enter the command in uppercase letters, lowercase letters, or any combination. In this example, you can enter CMDNA, CMDNAM, or CMDNAME in any combination of uppercase and lowercase letters.

```
>>--CMDNAme-----><
```

Note: Command names in descriptive text are always capitalized.

Required parameters

When a parameter is on the same line as the command name, the parameter is required. When two or more parameter values are in a stack and one of them is on the line, you *must* specify one value.

In this example, you must enter PARMNAME=A, PARMNAME=B, or PARMNAME=C. Do not include any blanks immediately before or after the equal sign (=).

```
>>-PARMName-----+A+-----><
      +-B-+
      '-C-'
```

Optional parameters

When a parameter is below the line, the parameter is optional. In this example, you can enter PARMNAME=A or nothing at all. Do not include any blanks immediately before or after the equal sign (=).

```
>>-+-----+-----><
      '-PARMName-----A-'
```

When two or more parameter values are in a stack below the line, all of them are optional. In this example, you can enter PARMNAME=A, PARMNAME=B, PARMNAME=C, or nothing at all. Do not include any blanks immediately before or after the equal sign (=).

```
>>-+-----+-----><
      '-PARMName-----A-+-'
              +-B-+
              '-C-'
```

Defaults

Defaults are above the line. The system uses the default unless you override it. You can override the default by entering an option from the stack below the line.

In this example, PARMNAME=A is the default. You can also enter PARMNAME=A, PARMNAME=B, or PARMNAME=C. Do not include any blanks before or after the equal sign (=).

```
.-PARMName-----A-----
>>-+-----+-----><
      '-PARMName-----A-+-'
              +-B-+
              '-C-'
```

Variables

Highlighted lowercase items (like this) denote variables. In these examples, var_name represents variables::

```
>>-CMDName--var_name-----><

>>-+-----+-----><
      '-PARMname-----var_name-'
```

Special characters

You must code these symbols exactly as they appear in the syntax diagram.

- * Asterisk
- :
- Colon
- ,

,	Comma
=	Equal sign
-	Hyphen
()	Parentheses
.	Period

Repeating values

An arrow returning to the left means that the item can be repeated. A character within the arrow means that you must separate repeated items with that character.

```

      .- ,-----
      v |
>>---file_name+-----><

```

Repeatable choices

A stack of values followed by an arrow returning to the left means that you can select more than one value or, when permitted, repeat a single item. In this example, you can choose more than one value, with each name delimited with a comma. Do not include any blanks before or after the equal sign (=).

```

      .- ,-----
      v |
>>-PARMName-----+value1+-----><
                    +-value2-+
                    '-value3-'

```

Footnotes

Footnotes are enclosed in parentheses.

```

      .- ,-----
      v (1) |
>>-----file_name+-----><

```

Notes:

1. You can specify up to five file names.

Entering parameters

The order in which you enter parameters can be important. The following example shows a portion of the command for defining a copy storage pool:

```

>>-DEFine STGpool--pool_name--device_class_name----->
>>-POOLtype----Copy+-----+----->
                    '-DESCRIPTION----description-'
      .-REclaim----100-----
>>+-----+-----><
    '-REclaim----percent-'

```

The first two parameters in this command (*pool_name* and *device_class_name*) are required parameters. *pool_name* and *device_class_name* are also positional. That is, they must be entered in the order shown, immediately after the command name. The POOLTYPE parameter is a required keyword parameter. DESCRIPTION and RECLAIM, are optional keyword parameters.

Keyword parameters are identified by an equal sign that specifies a specific value or a variable. Keyword parameters must follow any positional parameters in a command.

The following command entries, in which the keyword parameters are ordered differently, are both acceptable:

```
define stgpool mycopypool mydeviceclass pooltype=copy description=engineering
reclaim=50
define stgpool mycopypool mydeviceclass description=engineering pooltype=copy
reclaim=50
```

The following example, in which one of the positional parameters follows a keyword parameter, is not acceptable:

```
define stgpool mycopypool pooltype=copy mydeviceclass description=engineering
reclaim=50
```

Syntax fragments

Some diagrams, because of their length, must display parts of the syntax with fragments. The fragment name appears between vertical bars in the diagram.

The expanded fragment appears in the diagram after all other parameters or at the bottom of the diagram. A heading with the fragment name identifies the expanded fragment. Commands appearing directly on the line are required.

In this example, the fragment is named "Fragment".

```
>>-| Fragment |-----><
Fragment
.-A-.
|--+-----|
+-B-+
'-C-'
```

Using continuation characters to enter long commands

Continuation characters are useful when you want to process a command that is longer than your screen or window width. You can use continuation characters in the interactive mode of the administrative client.

About this task

Without continuation characters, you can enter up to 256 characters. With continuation characters, you can enter up to 1500 characters.

Note: In the MACRO command, the maximums apply after any substitution variables have been applied.

With continuation characters, you can do the following:

- Enter a dash at the end of the line you want to continue.

For example:

```
register admin pease mypasswd -
contact="david, ext1234"
```

- Continue a list of values by entering a dash or a back slash, with no preceding blank spaces, after the last comma of the list that you enter on the first line. Then, enter the remaining items in the list on the next line with no preceding blank spaces.

For example:

```
stgpools=stg1, stg2, stg3, -
stg4, stg5, stg6
```

- Continue a string of values that are enclosed in quotation marks by entering the first part of the string that is enclosed in quotation marks, followed by a dash or a back slash at the end of the line. Then, enter the remainder of the string on the next line, enclosed in the same type of quotation marks.

For example:


```
contact="david pease, bldg. 100, room 2b, san jose,"-  
"ext. 1234, alternate contact-norm pass,ext 2345"
```

IBM Spectrum Protect™ concatenates the two strings with no intervening blanks. You must use only this method to continue a quoted string of values across more than one line.

Naming IBM Spectrum Protect objects

IBM Spectrum Protect™ restricts the number and type of characters that you can use to name objects.

About this task

The following characters are available for defining object names.

Character	Description
A–Z	Any letter, A through Z
0–9	Any number, 0 through 9
_	Underscore
.	Period
-	Hyphen
+	Plus
&	Amperсанд

The following table shows the maximum length of characters permitted for naming objects.

Type of Name	Maximum Length
Administrators, client option sets, client nodes, passwords, server groups, server, names, virtual file space names	64
Restartable export identifiers	64
High-level and low-level TCP/IP (IPv4 or IPv6) addresses	64
Device classes, drives, libraries, management classes, policy domains, profiles, schedules scripts, backup sets, storage pools	30

The following characters are available for defining password names:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )  
| { } [ ] : ; < > , ? / ~
```

Passwords considered "LOCAL" are those passwords that authenticate with the IBM Spectrum Protect server and are not case-sensitive. Once a node or administrator is updated to use the SESSIONSECURITY=STRICT parameter, the password becomes case-sensitive the next time you change the it. Passwords considered "LDAP" are those passwords that authenticate with an LDAP directory server and are case-sensitive.

When you use DEFINE commands to define database, recovery log, and storage pool volumes, the naming convention for the volume name is dependent on the type of sequential access media or random access media that you are using. Refer to the specific VOLUME command for details.

Using wildcard characters to specify object names

In some commands, such as the query commands, you can use wildcard characters to create a pattern-matching expression that specifies more than one object. Using wildcard characters makes it easier to tailor a command to your needs.

About this task

The wildcard characters you use depend on the operating system from which you issue commands. For example, you can use wildcard characters such as an asterisk (*) to match any (0 or more) characters, or you can use a question mark (?) or a percent sign (%) to match exactly one character.

Table 1 provides references to wildcard characters for some operating systems. Use wildcard characters appropriate for your system.

Table 1. Wildcard characters by operating system

Operating system	Match any	Match exactly one
AIX®, Linux, Windows	*	?
TSO	*	%

For example, if you want to query all the management classes whose names begin with DEV in all the policy sets in DOMAIN1, and your system uses an asterisk as the *match-any* character, you can enter:

```
query mgmtclass domain1 * dev*
```

If your system uses a question mark as the *match-exactly-one* character, and you want to query the management classes in POLICYSET1 in DOMAIN1, you can enter:

```
query mgmtclass domain1 policyset1 mc?
```

IBM Spectrum Protect™ displays information about management classes with names MC.

Table 2 shows additional examples of using wildcard characters to match any characters.

Table 2. Match-any character

Pattern	Matches	Does not match
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers

Table 3 shows additional examples of using wildcard characters to match exactly one character. The question mark (?) can be replaced by a percent sign (%) if your platform uses that character instead of (?).

Table 3. Match-exactly-one character

Pattern	Matches	Does not match
ab?	abc	ab, abab, abzzzz
ab?rs	abfrs	abrs, abllrs
ab?ef?rs	abdefjrs	abefrs, abdefrs, abefjrs
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkrs

Specifying descriptions in keyword parameters

If a description (a string of text) for a parameter begins with a single or double quotation mark, or contains any embedded blanks or equal signs, you must surround the value with either single (') or double (") quotation marks.

About this task

The opening and closing quotation marks must be the same type of quotation marks. For example, if the opening quotation is a single quotation mark, the closing quotation mark must also be a single quotation mark.

For example, to register a new client node named Louie, with a password of secret, and with his title included as contact information, enter:

```
register node louie secret contact="manager of dept. 61f"
```

The following table presents ways of entering a description for the CONTACT parameter. The value can contain quotation marks, embedded blanks, or equal signs.

For this description	Enter this
----------------------	------------

For this description	Enter this
manager	contact=manager
manager's	contact="manager's" or contact='manager's'
"manager"	contact=""manager"" or contact=""manager""
manager's report	contact="manager's report" or contact='manager's report'
manager's "report"	contact='manager's "report"'
manager=dept. 61f	contact='manager=dept. 61f'
manager reports to dept. 61f	contact='manager reports to dept. 61f' or contact="manager reports to dept. 61f"

Controlling command processing

You can run some IBM Spectrum Protect™ commands sequentially or concurrently with other commands. You can also route commands from one server to other servers for processing.

About this task

- Server command processing
IBM Spectrum Protect processes administrator commands either in the foreground or in the background. Commands that process in the foreground must complete before you can issue another command. When commands are processing in the background, you can issue additional commands at any time.
- Stopping background processes
Use the CANCEL PROCESS command to cancel commands that generate background processes.

Server command processing

IBM Spectrum Protect™ processes administrator commands either in the foreground or in the background. Commands that process in the foreground must complete before you can issue another command. When commands are processing in the background, you can issue additional commands at any time.

Most IBM Spectrum Protect commands process in the foreground. For some commands that normally process in the background (for example, BACKUP DB), you can specify the WAIT parameter (WAIT=YES) with the command so that the command processes in the foreground. You might want to process a command in the foreground rather than in the background for any of the following reasons:

- To quickly determine whether a command completed successfully. When you issue a command that processes in the foreground, IBM Spectrum Protect sends a confirmation message that indicates that the command completed successfully. If you process the command in the background, you need to open operational reporting or query the activity log to determine whether the command completed successfully.
- To monitor server activities (for example, messages) on the administrative client as a command is being processed. This might be preferable to searching a long activity log after the command has completed.
- To be able to start another process immediately after a command completed. For example, you might specify WAIT=YES for a command that takes a short time to process so that, when the processing completes, you can immediately start processing another command.
- To serialize commands in an administrative script when it is important that one command completes before another begins.

Check the individual command description to determine whether a command has a WAIT parameter.

You can cancel commands that are processed in the foreground from the server console or from another administrative client session.

Each background process is assigned a process number. Use the QUERY PROCESS command to obtain the status and process number of a background process.

Note:

- If you are defining a schedule with a command that specifies WAIT=NO (the default), and you issue QUERY EVENT to determine the status of your scheduled operation, failed operations report an event status of COMPLETED with a return of

OK. In order for the QUERY EVENT output to reflect the failed status, the WAIT parameter must be set to YES. This runs the scheduled operation in the foreground and informs you of the status when it completes.

- You cannot process commands in the foreground from the server console.

Stopping background processes

Use the CANCEL PROCESS command to cancel commands that generate background processes.

About this task

Use the QUERY PROCESS command to obtain the status and process number of a background process. If a background process is active when you cancel it, the server stops the process. Any changes that are uncommitted are rolled back. However, changes that are committed are not rolled back.

When you issue a QUERY command from the administrative client, multiple screens of output might be generated. If this occurs and additional output is not needed, you can cancel the display of output to the client workstation. Doing so does not end the processing of the command.

Performing tasks concurrently on multiple servers

Command routing allows you to route commands to one or more servers for processing and then collect the output from these servers.

About this task

To route commands to other servers, you must have the same administrator ID and password as well as the required administrative authority on each server to which the command is being routed. You cannot route commands to other servers from the server console.

After the command has completed processing on all servers, the output displays, in its entirety, for each server. For example, the output from SERVER_A displays in its entirety, followed by the output from SERVER_B. The output includes summary messages for each individual server and identifies which server processed the output. Return codes indicate whether commands processed on the servers successfully. These return codes include one of three severities: 0, ERROR, or WARNING.

Each server that is identified as the target of a routed command must first be defined using the DEFINE SERVER command. The command is automatically routed to all servers specified as members of a server group or to individual servers specified with the command.

The following examples describe how to route the QUERY STGPOOL command to one server, multiple servers, a server group, multiple server groups, or a combination of servers and server groups. Each server or server group in a list must be separated with a comma, without spaces.

Routing commands to a single server

Procedure

To route the QUERY STGPOOL command to a server named ASTRO, enter:

```
astro: query stgpool
```

The colon after the server name indicates the end of the routing information. This is also called the *server prefix*. Another way to indicate the end of routing information is to use parentheses around the server name, for example:

```
(astro) query stgpool
```

Routing commands to multiple servers

About this task

Procedure

To route the QUERY STGPOOL command to multiple servers named HD_QTR, MIDAS, SATURN, enter:

```
hd_qtr,midas,saturn: query stgpool
```

If the first server has not been defined to IBM Spectrum Protect, the command is routed to the next defined server in the list of servers.

You can also enter the command this way:

```
(hd_qtr,midas,saturn) query stgpool
```

Routing commands to a server group

About this task

In this example, the server group ADMIN has servers named SECURITY, PAYROLL, PERSONNEL defined as group members. The command is routed to each of these servers.

Procedure

To route the QUERY STGPOOL command to the server group named ADMIN, enter:

```
admin: query stgpool
```

You can also enter the command this way:

```
(admin) query stgpool
```

Routing commands to server groups

About this task

In this example, the server group ADMIN2 has servers SERVER_A, SERVER_B, and SERVER_C defined as group members, and server group ADMIN3 has servers ASTRO, GUMBY, and CRUSTY defined as group members. The command is routed to servers SERVER_A, SERVER_B, SERVER_C, ASTRO, GUMBY, and CRUSTY.

Procedure

To route the QUERY STGPOOL command to two server groups that are named ADMIN2 and ADMIN3, enter:

```
admin2,admin3: query stgpool
```

You can also enter the command this way:

```
(admin2,admin3) query stgpool
```

Routing commands to two servers and a server group

About this task

In this example, the server group DEV_GROUP has servers SALES, MARKETING, and STAFF defined as group members. The command is routed to servers SALES, MARKETING, STAFF, MERCURY, and JUPITER.

Procedure

To route the QUERY STGPOOL command to a server group named DEV_GROUP and to the servers named MERCURY and JUPITER, enter:

```
dev_group,mercury,jupiter: query stgpool
```

You can also enter the command this way:

```
(dev_group,mercury,jupiter) query stgpool
```

Routing commands inside scripts

About this task

When routing commands inside scripts, you must enclose the server or server group in parentheses and omit the colon. Otherwise, the command will not be routed when the RUN command is issued, and will only be run on the server where the RUN command is issued.

For example, to route the QUERY STGPOOL command inside a script:

Procedure

1. Define a script called QU_STG to route it to the DEV_GROUP server group.

```
define script qu_stg "(dev_group) query stgpool"
```

2. Run the QU_STG script:

```
run qu_stg
```

Results

In this example, the server group DEV_GROUP has servers SALES, MARKETING, and STAFF defined as group members. The QUERY STGPOOL command is routed to these servers.

Privilege classes for commands

The authority granted to an administrator through the privilege class determines which administrative commands that the administrator can issue.

There are four administrator privilege classes in IBM Spectrum Protect™:

- System
- Policy
- Storage
- Operator

After an administrator has been registered using the REGISTER ADMIN command, the administrator can issue a limited set of commands, including all query commands. When you install IBM Spectrum Protect, the server console is defined as a system administrator named SERVER_CONSOLE and is granted system privilege.

- **Commands requiring system privilege**
An administrator with system privilege has the highest level of authority for the server. With system privilege, an administrator can issue any administrative command and has authority to manage all policy domains and all storage pools.
- **Commands requiring policy privilege**
An administrator with policy privilege can issue commands that relate to policy management objects such as policy domains, policy sets, management classes, copy groups, and schedules. The policy privilege can be unrestricted, or can be restricted to specific policy domains.
- **Commands requiring storage privilege**
An administrator with storage privilege can issue commands that allocate and control storage resources for the server. The storage privilege can be unrestricted, or can be restricted to specific storage pools.
- **Commands requiring operator privilege**
An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.
- **Commands any administrator can issue**
A limited number of commands can be used by any administrator, even if that administrator has not been granted any specific administrator privileges.

Commands requiring system privilege

An administrator with system privilege has the highest level of authority for the server. With system privilege, an administrator can issue any administrative command and has authority to manage all policy domains and all storage pools.

Table 1 lists the commands that administrators with system privilege can issue. In some cases administrators with lower levels of authority, for example, unrestricted storage privilege, can also issue these commands. In addition, the REQSYSAUTHOUTFILE server option can be used to specify that certain commands require system privilege if they cause the server to write to an external file. For more information about this server option, review REQSYSAUTHOUTFILE.

Table 1. System privilege commands

Command name	Command name
<ul style="list-style-type: none"> • AUDIT LDAPDIRECTORY • AUDIT LICENSES • ACCEPT DATE • BEGIN EVENTLOGGING • CANCEL EXPIRATION • CANCEL PROCESS • CANCEL REPLICATION • CANCEL REQUEST • CANCEL RESTORE • CLEAN DRIVE • COPY ACTIVATEDATA • COPY DOMAIN • COPY POLICYSET • COPY PROFILE • COPY SCHEDULE (Review note.) • COPY SCRIPT • COPY SERVERGROUP • DEFINE BACKUPSET • DEFINE CLIENTACTION • DEFINE CLIENTOPT • DEFINE CLOPTSET • DEFINE COLLOGGROUP • DEFINE COLLOCMEMBER • DEFINE DEVCLASS • DEFINE DOMAIN • DEFINE DRIVE • DEFINE EVENTSERVER • DEFINE GRPMEMBER • DEFINE LIBRARY • DEFINE MACHINE • DEFINE MACHNODEASSOCIATION • DEFINE NODEGROUP • DEFINE NODEGROUPMEMBER • DEFINE PATH • DEFINE PROFASSOCIATION • DEFINE PROFILE • DEFINE RECMEDMACHASSOCIATION • DEFINE RECOVERYMEDIA • DEFINE SCHEDULE (Review note.) • DEFINE SCRIPT • DEFINE SERVER • DEFINE SERVERGROUP 	<ul style="list-style-type: none"> • DEFINE SPACETRIGGER • DEFINE STGPOOL • DEFINE SUBSCRIPTION • DEFINE VIRTUALFSMAPPING • DEFINE VOLUME • DELETE BACKUPSET • DELETE CLIENTOPT • DELETE CLOPTSET • DEFINE COLLOGGROUP • DEFINE COLLOCMEMBER • DELETE DOMAIN • DELETE DRIVE • DELETE EVENTSERVER • DELETE GRPMEMBER • DELETE LIBRARY • DELETE MACHINE • DELETE MACHNODEASSOCIATION • DELETE NODEGROUP • DELETE NODEGROUPMEMBER • DELETE PROFASSOCIATION • DELETE PROFILE • DELETE RECMEDMACHASSOCIATION • DELETE RECOVERYMEDIA • DELETE SCHEDULE (Review note.) • DELETE SCRIPT • DELETE SERVER • DELETE SERVERGROUP • DELETE SPACETRIGGER • DELETE STGPOOL • DELETE SUBSCRIBER • DELETE SUBSCRIPTION • DELETE VIRTUALFSMAPPING • DISABLE EVENTS • ENABLE EVENTS • END EVENTLOGGING • EXPIRE INVENTORY • EXPORT ADMIN • EXPORT NODE • EXPORT POLICY • EXPORT SERVER • GENERATE BACKUPSET • GRANT AUTHORITY

Command name	Command name
<ul style="list-style-type: none"> • GRANT PROXYNODE • IDENTIFY DUPLICATES • IMPORT NODE • IMPORT POLICY • IMPORT SERVER • INSERT MACHINE • LABEL LIBVOLUME • LOCK ADMIN • LOCK PROFILE • MIGRATE STGPOOL • MOVE DRMEDIA • MOVE MEDIA • MOVE GRPMEMBER • NOTIFY SUBSCRIBERS • PERFORM LIBACTION • PING SERVER • PREPARE • QUERY BACKUPSETCONTENTS • QUERY MEDIA • QUERY RPFCONTENT • QUERY TOC • RECLAIM STGPOOL • RECONCILE VOLUMES • REGISTER ADMIN • REGISTER LICENSE • REMOVE ADMIN • REMOVE REPLNODE • RENAME ADMIN • RENAME SCRIPT • RENAME SERVERGROUP • RENAME STGPOOL • REPLICATE NODE • RESET PASSEXP • RESTORE NODE • REVOKE AUTHORITY • REVOKE PROXYNODE • RUN • SET ACCOUNTING • SET ACTLOGRETENTION • SET ARCHIVERETENTIONPROTECTION • SET ARREPLRULEDEFAULT • SET BKREPLRULEDEFAULT • SET CLIENTACTDURATION 	<ul style="list-style-type: none"> • SET CONFIGMANAGER • SET CONFIGREFRESH • SET CONTEXTMESSAGING • SET CROSSDEFINE • SET DBRECOVERY • SET DEFAULTAUTHENTICATION • SET DRMACTIVEDATASTGPOOL • SET DRMCHECKLABEL • SET DRMCMDFILENAME • SET DRMCOPYCONTAINERSTGPOOL • SET DRMCOPYSTGPOOL • SET DRMCOURIERNAME • SET DRMDBBACKUPEXPIREDAYS • SET DRMFILEPROCESS • SET DRMINSTRPREFIX • SET DRMNOTMOUNTABLENAME • SET DRMPPLANPREFIX • SET DRMPPLANVPOSTFIX • SET DRMPRIMSTGPOOL • SET DRMRPFEXPIREDAYS • SET DRMVaultNAME • SET EVENTRETENTION • SET INVALIDPWLIMIT • SET LDAPPASSWORD • SET LDAPUSER • SET LICENSEAUDITPERIOD • SET MAXCMDRETRIES • SET MAXSCHEDSESSIONS • SET MINPWLENGTH • SET PASSEXP • SET QUERYSCHEDPERIOD • SET RANDOMIZE • SET REPLRETENTION • SET REPLSERVER • SET RETRYPERIOD • SET SCHEDMODES • SET SERVERHLADDRESS • SET SERVERLLADDRESS • SET SERVERNAME • SET SERVERPASSWORD • SET SPREPLRULEDEFAULT • SET SUBFILE • SET TOCLOADRETENTION
<ul style="list-style-type: none"> • SETOPT • UNLOCK ADMIN • UNLOCK PROFILE • UPDATE ADMIN • UPDATE BACKUPSET • UPDATE CLIENTOPT • UPDATE CLOPTSET • UPDATE COLLOGGROUP • UPDATE DEVCLASS • UPDATE DRIVE • UPDATE LIBRARY • UPDATE LIBVOLUME • UPDATE MACHINE 	<ul style="list-style-type: none"> • UPDATE NODEGROUP • UPDATE PATH • UPDATE PROFILE • UPDATE RECOVERYMEDIA • UPDATE REPLRULE • UPDATE SCHEDULE (Review note.) • UPDATE SCRIPT • UPDATE SERVER • UPDATE SERVERGROUP • UPDATE SPACETRIGGER • UPDATE VIRTUALFSMAPPING • UPDATE VOLHISTORY • VALIDATE LANFREE • VALIDATE REPLICATION

Command name	Command name
Note: This command is restricted by the authority that is granted to an administrator. System privilege is required only for administrative command schedules. System or policy privilege is required for client operation schedules.	

Commands requiring policy privilege

An administrator with policy privilege can issue commands that relate to policy management objects such as policy domains, policy sets, management classes, copy groups, and schedules. The policy privilege can be unrestricted, or can be restricted to specific policy domains.

With unrestricted policy privilege, you can issue all of the administrator commands that require policy privilege. You can issue commands that affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains.

With restricted policy privilege, you can issue administrator commands that affect one or more policy domains for which authority is granted. For example, the DELETE MGMTCLASS command requires you to have policy privilege for the policy domain to which the management class belongs.

Table 1 lists the commands that an administrator with policy privilege can issue.

Table 1. Policy privilege commands

Command name	Command name
<ul style="list-style-type: none"> • ACTIVATE POLICYSET • ASSIGN DEFMGMTCLASS • CLEAN DRIVE • BACKUP NODE • COPY MGMTCLASS • COPY POLICYSET • COPY SCHEDULE (Review note 2.) • DEFINE ASSOCIATION • DEFINE BACKUPSET • DEFINE COPYGROUP • DEFINE CLIENTACTION • DEFINE CLIENTOPT • DEFINE MGMTCLASS • DEFINE NODEGROUP • DEFINE NODEGROUPMEMBER • DEFINE POLICYSET • DEFINE SCHEDULE • DELETE ASSOCIATION • DELETE BACKUPSET • DELETE COPYGROUP • DELETE EVENT (Review note 1.) • DELETE FILESPACE • DELETE MGMTCLASS • DELETE NODEGROUP • DELETE NODEGROUPMEMBER 	<ul style="list-style-type: none"> • DELETE POLICYSET • DELETE PATH • DELETE SCHEDULE (Review note 2.) • GENERATE BACKUPSET • LOCK NODE • QUERY BACKUPSETCONTENTS • REGISTER NODE • REMOVE NODE • RENAME FILESPACE • RENAME NODE • SET SUMMARYRETENTION • RESTORE NODE • QUERY TOC • UNLOCK NODE • UPDATE BACKUPSET • UPDATE COPYGROUP • UPDATE DOMAIN • UPDATE MGMTCLASS • UPDATE NODE • UPDATE NODEGROUP • UPDATE POLICYSET • UPDATE SCHEDULE (Review note 2.) • VALIDATE POLICYSET
Notes:	
<ol style="list-style-type: none"> 1. This command can be restricted by policy domain. An administrator with unrestricted policy privilege or restricted policy privilege for a specified policy domain can issue this command. 2. This command is restricted by the authority that is granted to an administrator. System privilege is required only for administrative command schedules. System or policy privilege is required for client operation schedules. 	

Commands requiring storage privilege

An administrator with storage privilege can issue commands that allocate and control storage resources for the server. The storage privilege can be unrestricted, or can be restricted to specific storage pools.

Unrestricted storage privilege permits you to issue all of the administrator commands that require storage privilege. You can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. You can also issue commands that affect the database and the recovery log. An unrestricted storage administrator cannot define or delete storage pools.

Restricted storage privilege permits you to issue administrator commands that only affect a storage pool for which you have been granted authority. For example, the DELETE VOLUME command only affects a storage pool volume that is defined to a specific storage pool.

Table 1 lists the commands an administrator with storage privilege can issue.

Table 1. Storage privilege commands

Command name	Command name
<ul style="list-style-type: none"> • AUDIT LIBRARY • AUDIT VOLUME (Review note.) • BACKUP DB • BACKUP DEVCONFIG • BACKUP STGPOOL • BACKUP VOLHISTORY • CHECKIN LIBVOLUME • CHECKOUT LIBVOLUME • COPY ACTIVATEDATA (Review note.) • DEFINE COLLOGGROUP • DEFINE COLLOCMEMBER • DEFINE DATAMOVER • DEFINE DEVCLASS • DEFINE DRIVE • DEFINE LIBRARY • DEFINE PATH • DEFINE VIRTUALFSMAPPING • DEFINE VOLUME (Review note.) • DEFINE SPACETRIGGER • DELETE COLLOGGROUP • DELETE COLLOCMEMBER • DELETE DATAMOVER • DELETE DEVCLASS • DELETE DRIVE • DELETE LIBRARY • DELETE PATH 	<ul style="list-style-type: none"> • DELETE SPACETRIGGER • DELETE VIRTUALFSMAPPING • DELETE VOLHISTORY • DELETE VOLUME (Review note.) • GRANT PROXYNODE • LABEL LIBVOLUME • MIGRATE STGPOOL • MOVE DATA (Review note.) • MOVE MEDIA • QUERY TAPEALERTMSG • RECLAIM STGPOOL • RESTORE STGPOOL • RESTORE VOLUME • REVOKE PROXYNODE • SET TAPEALERTMSG • UPDATE COLLOGGROUP • UPDATE DATAMOVER • UPDATE DEVCLASS • UPDATE DRIVE • UPDATE LIBRARY • UPDATE PATH • UPDATE SPACETRIGGER • UPDATE STGPOOL (Review note.) • UPDATE VIRTUALFSMAPPING
<p>Note: This command can be restricted by storage pool. An administrator with unrestricted storage privilege or restricted storage privilege for a specified storage pool can issue this command.</p>	

Commands requiring operator privilege

An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Table 1 lists the commands an administrator with operator privilege can issue.

Table 1. Operator privilege commands

Command Name	Command Name
--------------	--------------

Command Name	Command Name
<ul style="list-style-type: none"> • CANCEL SESSION • DISABLE SESSIONS • DISMOUNT VOLUME • ENABLE SESSIONS • HALT 	<ul style="list-style-type: none"> • MOVE DRMEDIA • MOVE MEDIA • QUERY MEDIA • REPLY • UPDATE VOLUME • VARY

Commands any administrator can issue

A limited number of commands can be used by any administrator, even if that administrator has not been granted any specific administrator privileges.

Table 1 lists the commands any registered administrator can issue.

Table 1. Commands issued by all administrators

Command Name	Command Name
<ul style="list-style-type: none"> • COMMIT • HELP • ISSUE MESSAGE • MACRO • PARALLEL • QUERY ACTLOG • QUERY ADMIN • QUERY ASSOCIATION • QUERY AUDITOCUPANCY • QUERY BACKUPSET • QUERY CLOPTSET • QUERY COLLOGROUP • QUERY CONTENT • QUERY COPYGROUP • QUERY DATAMOVER • QUERY DB • QUERY DBSPACE • QUERY DEVCLASS • QUERY DIRSPACE • QUERY DOMAIN • QUERY DRIVE • QUERY DRMEDIA • QUERY DRMSTATUS • QUERY ENABLED • QUERY EVENT • QUERY EVENTRULES • QUERY EVENTSERVER • QUERY FILESPACE • QUERY LIBRARY • QUERY LIBVOLUME • QUERY LICENSE • QUERY LOG • QUERY MACHINE • QUERY MGMTCLASS • QUERY MOUNT • QUERY NASBACKUP 	<ul style="list-style-type: none"> • QUERY NODE • QUERY NODEDATA • QUERY NODEGROUP • QUERY OCCUPANCY • QUERY OPTION • QUERY PATH • QUERY POLICYSET • QUERY PROCESS • QUERY PROFILE • QUERY PROXYNODE • QUERY RECOVERYMEDIA • QUERY REPLICATION • QUERY REPLNODE • QUERY REPLRULE • QUERY REQUEST • QUERY RESTORE • QUERY RPFIL • QUERY SCHEDULE • QUERY SCRIPT • QUERY SERVER • QUERY SERVERGROUP • QUERY SESSION • QUERY SPACETRIGGER • QUERY STATUS • QUERY STGPOOL • QUERY SUBSCRIBER • QUERY SUBSCRIPTION • QUERY SYSTEM • QUERY VIRTUALFSMAPPING • QUERY VOLHISTORY • QUERY VOLUME • QUIT • ROLLBACK • SELECT • SERIAL

Administrative commands

Administrative commands are available to manage and configure the server.

Information for each command includes:

- A description of the tasks a command performs
 - The administrator privilege class required to use the command
 - A syntax diagram that identifies the required and optional parameters for the command
 - Descriptions of each parameter of the command
 - Examples of using the command
 - A list of related commands
-
- **ACCEPT DATE** (Accepts the current system date)
Use this command to allow the server to begin normal processing, when the server does not start normal processing because of a discrepancy between the server date and the current date on the system.
 - **ACTIVATE POLICYSET** (Activate a new policy set)
Use this command to copy the contents of a policy set to the ACTIVE policy set for the domain. The server uses the rules in the ACTIVE policy set to manage client operations in the domain. You can define multiple policy sets for a policy domain, but only one policy set can be active. The current ACTIVE policy set is replaced by the one you specify when you issue this command. You can modify the ACTIVE policy set only by activating another policy set.
 - **ASSIGN DEFMGMTCLASS** (Assign a default management class)
Use this command to specify a management class as the default management class for a policy set. You must assign a default management class for a policy set before you can activate that policy set.
 - **AUDIT** commands
Use the AUDIT commands to review or examine the adequacy of the database information and the storage pool volume. The AUDIT LDAPDIRECTORY command deletes nodes or administrator IDs from an LDAP directory server, that do not authenticate their passwords with the LDAP directory server.
 - **BACKUP** commands
Use the BACKUP commands to create backup copies of IBM Spectrum Protect™ information or objects.
 - **BEGIN EVENTLOGGING** (Begin logging events)
Use this command to begin logging events to one or more receivers. A receiver for which event logging has begun is an *active receiver*.
 - **CANCEL** commands
Use the CANCEL commands to end a task or process before it is completed.
 - **CHECKIN LIBVOLUME** (Check a storage volume into a library)
Use this command to add a sequential access storage volume or a cleaning tape to the server inventory for an automated library. The server cannot use a volume that physically resides in an automated library until that volume is checked in.
 - **CHECKOUT LIBVOLUME** (Check a storage volume out of a library)
Use this command to remove a sequential access storage volume from the server inventory for an automated library. This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.
 - **CLEAN DRIVE** (Clean a drive)
Use this command when you want IBM Spectrum Protect to immediately load a cleaner cartridge into a drive regardless of the cleaning frequency.
 - **COMMIT** (Control committing of commands in a macro)
Use this command to control when a command is committed in a macro and to update the database when commands complete processing. When issued from the console mode of the administrative client, this command does not generate a message.
 - **CONVERT STGPOOL** (Convert a storage pool to a container storage pool)
Use this command to convert a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL) to a directory-container or a cloud-container storage pool. You can use container storage pools for both inline and client-side data deduplication.
 - **COPY** commands
Use the COPY commands to create a copy of IBM Spectrum Protect objects or data.
 - **DEACTIVATE DATA** (Deactivate data for a client node)
Use this command to specify that active data that was backed up for an application client node before a specified date is no longer needed. The command marks the data as inactive so it can be deleted according to your data retention policies.
 - **DECOMMISSION** commands
Use the DECOMMISSION commands to remove client nodes from the production environment. Client nodes include applications, systems, and virtual machines.
 - **DEFINE** commands
Use the DEFINE commands to create IBM Spectrum Protect objects.
 - **DELETE** commands
Use the DELETE commands to delete or remove an IBM Spectrum Protect object.

- **DISABLE commands**
Use DISABLE commands to prevent some types of operations by the server.
- **DISMOUNT command**
Use the DISMOUNT command to dismount a volume by the real device address or by volume name.
- **DISPLAY OBJNAME (Display a full object name)**
Use this command when you want IBM Spectrum Protect to display a full object name if the name displayed in a message or query output has been abbreviated due to length. Object names that are very long can be difficult to display and use through normal operating system facilities. The IBM Spectrum Protect server will abbreviate long names and assign them a token ID which might be used if the object path name exceeds 1024 bytes. The token ID is displayed in a string that includes identifiers for the node, filespace, and object name. The format is: [TSMOBJ:nID.fsID.objID]. When specified with the DISPLAY OBJNAME command, the token ID can be used to show the full object name.
- **ENABLE commands**
Use ENABLE commands to allow some types of operations by the server.
- **ENCRYPT STGPOOL (Encrypt data in a storage pool)**
Use this command to encrypt data in a directory-container or cloud-container storage pool.
- **END EVENTLOGGING (Stop logging events)**
Use this command to stop logging events to an active receiver.
- **EXPIRE INVENTORY (Manually start inventory expiration processing)**
Use this command to manually start inventory expiration processing. The inventory expiration process removes client backup and archive file copies from server storage. Removal is based on policy specifications in the backup and archive copy groups of the management classes to which the files are bound.
- **EXPORT commands**
Use the EXPORT commands to copy information from an IBM Spectrum Protect server to sequential removable media.
- **EXTEND DBSPACE (Increase space for the database)**
Use this command to increase space for the database by adding directories for the database to use.
- **GENERATE commands**
Use the GENERATE commands for backup sets for a selected filespace or client node.
- **GRANT commands**
Use the GRANT command to grant appropriate privileges or access.
- **HALT (Shut down the server)**
Use this command to shut down the server. The HALT command forces an abrupt shutdown, which cancels all the administrative and client node sessions even if they are not completed.
- **HELP (Get help on commands and error messages)**
Use this command to display administrative commands and error messages. You can issue the command from an administrative command line client.
- **IDENTIFY DUPLICATES (Identify duplicate data in a storage pool)**
Use this command to start or stop processes that identify duplicate data in a storage pool. You can specify the number of duplicate-identification processes and their duration.
- **IMPORT commands**
Use the IMPORT commands to import information from export media to an IBM Spectrum Protect server.
- **INSERT MACHINE (Insert machine characteristics information or recovery instructions)**
Use this command to add client machine characteristics or recovery instructions to existing machine information in the database.
- **ISSUE MESSAGE (Issue a message from a server script)**
Use this command with return code processing in a script to issue a message from a server script to determine where the problem is with a command in the script.
- **LABEL LIBVOLUME (Label a library volume)**
Use this command to label tape volumes or, in an automated library, to label the volumes automatically as they are checked in. With this command, the server uses the full-length label with which the volumes are often pre-labeled.
- **LOAD DEFALERTTRIGGERS (Load the default set of alert triggers)**
Use this command to load the default set of alert triggers to the IBM Spectrum Protect server.
- **LOCK commands**
Use the LOCK command to prevent users from accessing the server.
- **MACRO (Invoke a macro)**
Use this command to invoke a file from the administrative command line that contains one or more IBM Spectrum Protect administrative commands to be performed.
- **MIGRATE STGPOOL (Migrate storage pool to next storage pool)**
Use this command to migrate files from one storage pool to the next storage pool in the storage hierarchy.
- **MOVE commands**
Use the MOVE commands to either transfer backup or archive data between storage pools, or to move disaster recovery media on and off site.

- NOTIFY SUBSCRIBERS (Notify managed servers to update profiles)
Use this command on a configuration manager to notify one or more managed servers to request that their configuration information be immediately refreshed.
- PERFORM LIBACTION (Define or delete all drives and paths for a library)
Use this command to define or delete all drives and their paths for a single library in one step.
- PING SERVER (Test the connection between servers)
Use this command to test the connection between the local server and a remote server.
- PREPARE (Create a recovery plan file)
Use this command to create a recovery plan file, which contains the information that is needed to recover an IBM Spectrum Protect server. You can store a recovery plan file on a file system that is accessible to the source server or on a target server.
- PROTECT STGPOOL (Protect data that belongs to a storage pool)
Use this command to protect data in a directory-container storage pool by storing a copy of the data in another storage pool on a replication target server or on the same server by protecting the data to tape. When you protect the directory-container storage pool, you can later try to repair damage in the storage pool by using the REPAIR STGPOOL command.
- QUERY commands
Use the QUERY commands to request or display information about IBM Spectrum Protect objects.
- QUIT (End the interactive mode of the administrative client)
Use this command to end an administrative client session in interactive mode.
- RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)
Use this command to reclaim volumes in a sequential-access storage pool. Reclamation does not move inactive versions of backup data from volumes in active-data pools.
- RECONCILE VOLUMES (Reconcile differences in the virtual volume definitions)
Issue this command from the source server to reconcile differences between virtual volume definitions on the source server and archive files on the target server. IBM Spectrum Protect finds all volumes of the specified device class on the source server and all corresponding archive files on the target server. The target server inventory is also compared to the local definition for virtual volumes to see if inconsistencies exist.
- REGISTER commands
Use the REGISTER commands to define or add objects to IBM Spectrum Protect.
- REMOVE commands
Use the REMOVE commands to remove an object from IBM Spectrum Protect.
- RENAME commands
Use the RENAME commands to change the name of an existing object.
- REPAIR STGPOOL (Repair a directory-container storage pool)
Use this command to repair deduplicated extents in a directory-container storage pool. Damaged deduplicated extents are repaired with extents that are backed up to the target replication server or to container-copy storage pools on the same server.
- REPLICATE NODE (Replicate data in file spaces that belong to a client node)
Use this command to replicate data in file spaces that belong to one or more client nodes or defined groups of client nodes.
- REPLY (Allow a request to continue processing)
Use this command and an identification number to inform the server that you have completed a requested operation. Not all server requests require a reply. This command is required only if the request message specifically indicates that a reply is needed.
- RESET PASSEXP (Reset password expiration)
Use the RESET PASSEXP command to reset the password expiration period to the common expiration period for administrator and client node passwords. The RESET PASSEXP command does not apply to passwords that are stored on an LDAP directory server.
- RESTART EXPORT (Restart a suspended export operation)
Use this command to restart a suspended export operation.
- RESTORE commands
Use the RESTORE commands to restore IBM Spectrum Protect storage pools or volumes.
- REVOKE commands
Use the REVOKE commands to revoke privileges or access.
- ROLLBACK (Rollback uncommitted changes in a macro)
Use this command within a macro to undo any processing changes made by commands run by the server but not yet committed to the database. A committed change is permanent and cannot be rolled back. The ROLLBACK command is useful for testing macros.
- RUN (Run an IBM Spectrum Protect script)
Use this command to run an IBM Spectrum Protect script. To issue this command on another server, the script being run must be defined on that server.
- SELECT (Perform an SQL query of the IBM Spectrum Protect database)
Use the SELECT command to create and format a customized query of the IBM Spectrum Protect database.

- **SET commands**
Use the SET commands to specify values that affect many different IBM Spectrum Protect operations.
- **SETOPT (Set a server option for dynamic update)**
You can use the SETOPT command to update most server options dynamically without stopping and restarting the server. For the DBDIAGLOGSIZE option, you must stop and start the server. A SETOPT command contained in a macro or a script cannot be rolled back.
- **SHRED DATA (Shred data)**
Use this command to manually start the process of shredding deleted sensitive data. Manual shredding is possible only if automatic shredding is disabled.
- **SUSPEND EXPORT (Suspend a currently running export operation)**
Use this command to suspend a currently running server-to-server export operation which has a FILEDATA value that is not NONE. The export operation that you want to suspend must be past the initialization phase to be eligible for suspension. The state of the export operation is saved. The operation can be restarted by issuing the RESTART EXPORT command.
- **UNLOCK commands**
Use the UNLOCK commands to reestablish access after an object was locked.
- **UPDATE commands**
Use the UPDATE command to modify one or more attributes of an existing IBM Spectrum Protect object.
- **VALIDATE commands**
Use the VALIDATE command to verify that an object is complete or valid for IBM Spectrum Protect.
- **VARY (Bring a random access volume online or offline)**
Use this command to make a random access storage pool volume online or offline to the server.

ACCEPT DATE (Accepts the current system date)

Use this command to allow the server to begin normal processing, when the server does not start normal processing because of a discrepancy between the server date and the current date on the system.

When the server does not start normal processing because of a discrepancy between the server date and the current date, this command forces the server to accept the current date and time as valid. If the system time is valid and the server has not been run for an extended time, this command should be run to allow the server to begin normal processing.

Attention: If the system date is invalid or the server was created or run previously with an invalid system date and this command is issued, any server processing or command that uses dates can have unexpected results. File expiration can be affected, for example. When the server is started with the correct date, files backed up with future dates will not be considered for expiration until that future date is reached. Files backed up with dates that have passed will expire faster. When the server processing encounters a future date, an error message is issued.

If the server detects an invalid date or time, server sessions become disabled (as if the DISABLE SESSIONS command had been issued). Expiration, migration, reclamation, and volume history deletion operations are not able to continue processing.

Use the ENABLE SESSIONS ALL command after you issue the ACCEPT DATE command to re-enable sessions to start.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-ACcEpt Date-----<<
```

Parameters

None.

Example: Accept the current system date

Allow the server to accept the current date as the valid date.

```
accept date
```

Related commands

Table 1. Command related to ACCEPT DATE

Command	Description
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.

ACTIVATE POLICYSET (Activate a new policy set)

Use this command to copy the contents of a policy set to the ACTIVE policy set for the domain. The server uses the rules in the ACTIVE policy set to manage client operations in the domain. You can define multiple policy sets for a policy domain, but only one policy set can be active. The current ACTIVE policy set is replaced by the one you specify when you issue this command. You can modify the ACTIVE policy set only by activating another policy set.

Before activating a policy set, check that the policy set is complete and valid by using the VALIDATE POLICYSET command.

The ACTIVATE POLICYSET command fails if any of the following conditions exist:

- A copy group specifies a copy storage pool as a destination.
- A management class specifies a copy storage pool as the destination for files that were migrated by an IBM Spectrum Protect™ for Space Management client.
- The policy set has no default management class.
- A TOCDESTINATION parameter is specified, and the storage pool is either a copy pool or has a data format other than NATIVE or NONBLOCK.

The ACTIVE policy set and the last activated policy set are not necessarily identical. You can modify the original policy set that you activated without affecting the ACTIVE policy set.

If the server has data retention protection enabled, the following conditions must exist:

- All management classes in the policy set to be activated must contain an archive copy group.
- If a management class exists in the active policy set, a management class with the same name must exist in the policy set to be activated.
- If an archive copy group exists in the active policy set, the corresponding copy group in the policy set to be activated must have a RETVER value at least as large as the corresponding values in the active copy group.

Attention: Retention protection only applies to archive objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-ACTivate Policyset--domain_name--policy_set_name-----><
```

Parameters

domain_name (Required)

Specifies the policy domain for which you want to activate a policy set.

policy_set_name (Required)

Specifies the policy set to activate.

Example: Activate a policy set on a specific policy domain

Activate the VACATION policy set in the EMPLOYEE_RECORDS policy domain.

```
activate policyset employee_records vacation
```


Related commands

Table 1. Commands related to ACTIVATE POLICYSET

Command	Description
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY DOMAIN	Displays information about policy domains.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

ASSIGN DEFMGMTCLASS (Assign a default management class)

Use this command to specify a management class as the default management class for a policy set. You must assign a default management class for a policy set before you can activate that policy set.

To ensure that clients can always back up and archive files, choose a default management class that contains both an archive copy group and a backup copy group.

The server uses the default management class to manage client files when a management class is not otherwise assigned or appropriate. For example, the server uses the default management class when a user does not specify a management class in the include-exclude list.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-ASsign DEFMGmtclass--domain_name--policy_set_name--class_name-><
```

Parameters

domain_name (Required)

Specifies the policy domain to which the management class belongs.

policy_set_name (Required)

Specifies the policy set for which you want to assign a default management class. You cannot assign a default management class to the ACTIVE policy set.

class_name (Required)

Specifies the management class that is to be the default management class for the policy set.

Example: Assign a default management class

Assign DEFAULT1 as the default management class for policy set SUMMER in the PROG1 policy domain.

```
assign defmgmtclass prog1 summer default1
```

Related commands

Table 1. Commands related to ASSIGN DEFMGMTCLASS

Command	Description
---------	-------------

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

AUDIT commands

Use the AUDIT commands to review or examine the adequacy of the database information and the storage pool volume. The AUDIT LDAPDIRECTORY command deletes nodes or administrator IDs from an LDAP directory server, that do not authenticate their passwords with the LDAP directory server.

- AUDIT CONTAINER
 - AUDIT CONTAINER (Verify the consistency of database information for a cloud container)
 - AUDIT CONTAINER (Verify the consistency of database information for a directory-container)
- AUDIT LDAPDIRECTORY (Audit an LDAP directory server)
- AUDIT LIBRARY (Audit volume inventories in an automated library)
- AUDIT LIBVOLUME (Verify database information for a tape volume)
- AUDIT LICENSES (Audit server storage usage)
- AUDIT VOLUME (Verify database information for a storage pool volume)

AIX

Linux

Windows

AUDIT CONTAINER commands

Use the AUDIT CONTAINER command to scan for inconsistencies between database information and a container in either a cloud or a directory storage pool.

- AUDIT CONTAINER (Verify the consistency of database information for a cloud container)
Use this command to scan for inconsistencies between database information and a container in a cloud-container storage pool. Cloud-container storage pools are not supported on Linux on System z®.
- AUDIT CONTAINER (Verify the consistency of database information for a directory-container)
Use this command to scan for inconsistencies between database information and a container in a directory-container storage pool.

AUDIT CONTAINER (Verify the consistency of database information for a cloud container)

Use this command to scan for inconsistencies between database information and a container in a cloud-container storage pool. Cloud-container storage pools are not supported on Linux on System z®.

You can use this command to complete the following actions for a container in a cloud-container storage pool:

- Scan the contents of a container to validate the integrity of the data extents

- Remove data from a container that is marked as *damaged*, such as when a file has references in the server database, but has missing or corrupted data in the cloud
- Mark an entire container as damaged
- Remove data that is marked as *orphaned*, such as when an object stored in the cloud does not have a reference in the server database

Privilege class

To use this command, you must have system privilege, or unrestricted storage privilege.

Syntax

```
>>-AUDit CONTainer--+--container_name-----+-->
                    +-STGpool---pool_name-----+
                    '-STGpool---pool_name--STGPOOLDIrectory---directory_name-'

.-Action---SCANAll-----
>--+-----+----->
'-Action---+SCANAll-----'
      +-REMOVEDamaged-+
      +-MARKDamaged---+
      '-SCANDamaged---'

.-FORCEOrphandbdel---No-----
>--+-----+----->
'-FORCEOrphandbdel---+No---+'
      '-Yes-'

.-MAXProcess---4-----.-Wait---No-----
>--+-----+-----+----->
'-MAXProcess---number-' '-Wait---+No---+'
      '-Yes-'

.-BEGINDate---before_first_audit-.
>--+-----+----->
'-BEGINDate---begin_date-----'

.-BEGINTime---00:00:00---
>--+-----+----->
'-BEGINTime---begin_time-'

.-ENDDate---after_last_audit-. .-ENDTime---23:59:59-.
>--+-----+-----+-----><
'-ENDDate---end_date-----' '-ENDTime---end_time-'
```

Parameters

container_name

Specifies the name of the container that you want to audit. If you do not specify this parameter, you must specify a cloud-container storage pool.

STGpool

Specifies the name of the cloud-container storage pool that you want to audit. This parameter is optional. If you specify only this parameter, all containers that are defined to the storage pool are audited. If you do not specify this parameter, you must specify a container.

STGPOOLDIrectory

Specifies the name of the cloud-container storage pool directory that you want to audit. This parameter is optional. Restriction: You must specify a storage pool that uses local storage.

Action

Specifies what action the server takes when a container in a cloud-container storage pool is audited. This parameter is optional. You can specify one of the following values:

SCANAll

Specifies that the server identifies database records that refer to data extents with inconsistencies. A check is done for data in the cloud-container storage pool that does not match data in the server database. This value is the default. The server marks the data extent as damaged in the database.

Tip: If you specify the ACTION=SCANALL parameter on an IBM® Cloud Object Storage storage pool that uses a vault with name indexing disabled, the audit operation scans the entire vault to identify orphaned extents in each container. In this situation, specify WAIT=YES if you want the audit operation to wait for the scan for orphaned extents to complete before it reports the audit as complete. This scan for orphaned extents occurs only if you do not specify a container name. If you specify a container that is in a vault with name indexing disabled, the audit operation does not scan for orphaned extents.

REMOVEDamaged

Specifies that the server removes any references to damaged extents from the server database. The damaged extents are also removed from the cloud-container storage pool if found. The server also removes any orphaned extents from the cloud-container storage pool, and removes the references to these orphaned extents from the database, as specified by the FORCEORPHANDBDEL parameter.

MARKDamaged

Specifies that the server explicitly marks all data extents in the container as damaged.

SCANDamaged

Specifies that the server checks only the existing damaged extents in the container.

Important: If no connection to the cloud exists, the ACTION=SCANALL and ACTION=SCANDAMAGED parameters do not run. However, the ACTION=MARKDAMAGED parameter runs as expected without a cloud connection, and the ACTION=REMOVEDAMAGED parameter marks any damaged data as orphaned. As soon as the connection to the cloud returns, the server deletes the orphaned extents.

State reset condition: If the audit does not detect an error with a data extent that is marked as damaged, the state of the data extent is reset. The data extent can then be used. This condition provides a means for resetting the state of damaged data extents if errors are caused by a correctable problem. The SCANALL and SCANDAMAGED options are the only options that reset a damaged extent if it is found not to be damaged.

FORCEOrphandbdel

Specifies that the server forces the deletion of orphaned extents from the server database, even if they are not deleted from the cloud-container storage pool. This parameter is optional. If you specify this parameter, you must also specify the ACTION=REMOVEDAMAGED parameter. The following options are available:

Yes

Specifies that the server deletes any orphaned extents from the server database, even if they are not deleted from the cloud-container storage pool.

No

Specifies that the server keeps the orphaned extents in the server database if they cannot be deleted from the cloud-container storage pool. This value is the default.

MAXProcess

Specifies the maximum number of parallel processes to use for checking a container in a cloud-container storage pool. This parameter is optional. Enter a value in the range 1 - 99. The default value is 4.

Restriction: The server ignores this parameter when you use MAXPROCESS with the ACTION=REMOVEDAMAGED parameter.

Wait

Specifies whether the audit or verification operation is completed in the foreground or background. This parameter is optional. The following options are available:

No

Specifies that the operation is completed in the background. You can continue with other tasks when the command is processing. Messages that are related to the background process are displayed in the activity log file or the server console, depending on where the messages are logged. This value is the default.

Yes

Specifies that the operation is completed in the foreground. It might take a long time to complete the operation. The operation must complete before you can continue with other tasks. Messages are displayed in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify the WAIT=YES parameter from the server console.

BEGINDate

Specifies the date range value at which auditing should start. Containers that were last audited within the specified date range are audited. If you specify a time but do not specify a beginning date, the current date is used. If you do not specify a beginning and end date, all containers are audited. The default is the date before the first audit was completed for the container. This parameter is optional.

You can specify the date to begin the audit in one of the following ways:

Value	Description	Example
-------	-------------	---------

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/2016
TODAY	The current date.	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 <i>or</i> -7. To audit all containers that were audited in the last week, specify BEGINDATE=TODAY-7 <i>or</i> BEGINDATE= -7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include containers that were audited a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include containers that were audited on the 10th day of the current month.

BEGINTime

Specifies the time range value at which auditing should start. Containers that were last audited within the specified time range are audited. If you do not specify a beginning and end time, the time range is set from 00:00:00 to 23:59:59. The default is 00:00:00. If you did not specify a date range, the default is today's date. This parameter is optional.

You can specify the date to begin the audit in one of the following ways:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date.	10:30:08
NOW	The current time on the specified begin date.	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified begin date.	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with BEGINTIME=NOW+3 <i>or</i> BEGINTIME=+3, containers with a last audit time of 12:00 or later on the begin date are audited.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified begin date.	NOW-04:00 <i>or</i> -04:00. If you issue this command at 9:00 with BEGINTime=NOW-3:30 <i>or</i> BEGINTime= -3:30, IBM Spectrum Protect™ audits containers with a last audit time of 5:30 or later on the begin date.

ENDDate

Specifies the date range value at which auditing should stop. Containers that were last audited within the specified date range are audited. If you specify a time but do not specify a value, the current date is used. If you do not specify a beginning and end date, all containers are audited. The default is the date after the last audit was completed for the container. This parameter is optional.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/2016
TODAY	The current date.	TODAY

Value	Description	Example
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To include containers that were audited up to yesterday, you can specify ENDDATE=TODAY-1 or ENDDATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include containers that were audited a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include containers that were audited on the 10th day of the current month.

ENDTime

Specifies the time range value at which auditing should stop. Containers that were last audited within the specified time range are audited. If you do not specify a beginning and end time, the time range is set to 00:00:00 to 23:59:59. The default is 23:59:59. This parameter is optional.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date.	10:30:08
NOW	The current time on the specified end date.	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date.	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, containers with a last audit time of 12:00 or earlier on the end date you specify are audited.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date.	NOW-03:30 or -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, containers with a last audit time of 5:30 or earlier on the end date you specify are audited.

Example: Audit a specific container in a cloud-container storage pool

Audit the 42-00000my000example000container000 container in a cloud-container storage pool.

```
audit container 42-00000my000example000container000 action=scanall
```

Example: Audit a cloud-container storage pool within a specific time frame

Audit a cloud-container storage pool that is named POOL3 and only include containers from yesterday between 9:30 and 12:30.

```
audit container stgpool=pool3 begindate=today-1  
begintime=09:30:00 endtime=12:30:00
```

Table 1. Commands related to AUDIT CONTAINER

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY CONTAINER	Displays information about a container.

Command	Description
QUERY DAMAGED	Displays information about damaged files.

AIX Linux Windows

AUDIT CONTAINER (Verify the consistency of database information for a directory-container)

Use this command to scan for inconsistencies between database information and a container in a directory-container storage pool.

You can use this command to complete the following actions for a container in a directory-container storage pool:

- Scan the contents of a container to validate the integrity of the data extents
- Remove damaged data from a container
- Mark an entire container as damaged

Privilege class

To issue this command, you must have system privilege, or unrestricted storage privilege.

Syntax

```
>>-AUDit CONTainer--+-container_name-----+-->
      +-STGpool---pool_name-----+
      '-STGpool---pool_name--STGPOOLDirectory---directory_name-'

.-Action---SCANAll-----.
>--+-----+----->
  '-Action---SCANAll---+'
      +-REMOVEDamaged+
      +-MARKDamaged---+
      '-SCANDamaged---'

.-MAXProcess---4-----.-Wait---No-----.
>--+-----+----->
  '-MAXProcess---number-' '-Wait---No---+'
      '-Yes-'

.-BEGINDate---before_first_audit-.
>--+-----+----->
  '-BEGINDate---begin_date-----'

.-BEGINTime---00:00:00---.
>--+-----+----->
  '-BEGINTime---begin_time-'

.-ENDDate---after_last_audit-. .-ENDTime---23:59:59-.
>--+-----+-----><
  '-ENDDate---end_date-----' '-ENDTime---end_time-'
```

Parameters

container_name

Specifies the name of the container that you want to audit. If you do not specify this parameter, you must specify a directory-container storage pool.

STGpool

Specifies the name of the directory-container storage pool that you want to audit. This parameter is optional. If you specify only this parameter, all containers that are defined to the storage pool are audited. If you do not specify this parameter, you must specify a container.

STGPOOLDirectory

Specifies the name of the container storage pool directory that you want to audit. This parameter is optional. If you specify this parameter, all containers that are defined to the container storage pool directory are audited. To specify this parameter,

you must also specify a storage pool.

Action

Specifies what action the server takes when a container in a directory-container storage pool is audited. This parameter is optional. You can specify one of the following values:

SCANAll

Specifies that the server identifies database records that refer to data extents with inconsistencies. This value is the default. The server marks the data extent as damaged in the database.

Tip: If you used the PROTECT STGPOOL command on a directory-container storage pool on the target server, you can repair the damaged data extent by using the REPAIR STGPOOL command.

REMOVEDamaged

Specifies that the server removes any files from the database that reference the damaged data extent.

MARKDamaged

Specifies that the server explicitly marks all data extents in the container as damaged.

SCANDamaged

Specifies that the server checks only the existing damaged extents in the container.

State reset condition: If the audit does not detect an error with a data extent that is marked as damaged, the state of the data extent is reset. The data extent can then be used. This condition provides a means for resetting the state of damaged data extents if errors are caused by a correctable problem. The SCANALL and SCANDAMAGED options are the only options that reset a damaged extent if it is found not to be damaged.

MAXProcess

Specifies the maximum number of parallel processes to use for checking a container in a directory-container storage pool. This parameter is optional. Enter a value in the range 1 - 99. The default value is 4.

Wait

Specifies whether the audit or verification operation is completed in the foreground or background. This parameter is optional. The following options are available:

No

Specifies that the operation is completed in the background. You can continue with other tasks when the command is processing. Messages that are related to the background process are displayed in the activity log file or the server console, depending on where the messages are logged. This is the default value.

Yes

Specifies that the operation is completed in the foreground. It might take a long time to complete the operation. The operation must complete before you can continue with other tasks. Messages are displayed in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify the WAIT=YES parameter from the server console.

BEGINDate

Specifies the date range value at which auditing should start. Containers that were last audited within the specified date range are audited. If you specify a time but do not specify a beginning date, the current date is used. If you do not specify a beginning and end date, all containers are audited. The default is the date before the first audit was completed for the container. This parameter is optional.

You can specify the date to begin the audit in one of the following ways:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/2016
TODAY	The current date.	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To audit all containers that were audited in the last week, specify BEGINDATE=TODAY-7 or BEGINDATE= -7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include containers that were audited a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include containers that were audited on the 10th day of the current month.

BEGINTime

Specifies the time range value at which auditing should start. Containers that were last audited within the specified time range are audited. If you do not specify a beginning and end time, the time range is set from 00:00:00 to 23:59:59. The default is 00:00:00. If you did not specify a date range, the default is today's date. This parameter is optional.

You can specify the date to begin the audit in one of the following ways:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date.	10:30:08
NOW	The current time on the specified begin date.	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified begin date.	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with BEGINTIME=NOW+3 or BEGINTIME=+3, containers with a last audit time of 12:00 or later on the begin date are audited.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified begin date.	NOW-04:00 <i>or</i> -04:00. If you issue this command at 9:00 with BEGINTime=NOW-3:30 or BEGINTime= -3:30, IBM Spectrum Protect™ audits containers with a last audit time of 5:30 or later on the begin date.

ENDDate

Specifies the date range value at which auditing should stop. Containers that were last audited within the specified date range are audited. If you specify a time but do not specify a value, the current date is used. If you do not specify a beginning and end date, all containers are audited. The default is the date after the last audit was completed for the container. This parameter is optional.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/2016
TODAY	The current date.	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 <i>or</i> -1. To include containers that were audited up to yesterday, you can specify ENDDATE=TODAY-1 or ENDDATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include containers that were audited a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include containers that were audited on the 10th day of the current month.

ENDTime

Specifies the time range value at which auditing should stop. Containers that were last audited within the specified time range are audited. If you do not specify a beginning and end time, the time range is set to 00:00:00 to 23:59:59. The

default is 23:59:59. This parameter is optional.
 You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date.	10:30:08
NOW	The current time on the specified end date.	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified end date.	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, containers with a last audit time of 12:00 or earlier on the end date you specify are audited.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified end date.	NOW-03:30 <i>or</i> -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, containers with a last audit time of 5:30 or earlier on the end date you specify are audited.

Example: Audit a specific storage pool container

Audit the 0000000000000721.dcf storage pool container.

```
audit container n:\ddcont2\07\0000000000000721.dcf action=scanall
```

Example: Remove damaged data from a directory-container storage pool

Audit a directory-container storage pool that is named NEWDEDUP and remove damaged files.

```
audit container stgpool=newdedup action=removedamaged
```

Example: Mark as damaged all of the data in a directory-container storage pool

Audit a directory-container storage pool that is named NEWDEDUP and mark all files as damaged.

```
audit container stgpool=newdedup maxprocess=2 action=markdamaged
```

Example: Audit a directory-container storage pool within a specific time frame

Audit a directory-container storage pool that is named POOL2 and only include containers before yesterday between 9:30 and 12:30.

```
audit container stgpool=pool2 begindate=today-1  
begintime=09:30:00 endtime=12:30:00
```

Table 1. Commands related to AUDIT CONTAINER

Command	Description
CANCEL PROCESS	Cancel a background server process.
MOVE CONTAINER	Moves the contents of a storage pool container to another container.
QUERY DAMAGED	Displays information about damaged files.

AUDIT LDAPDIRECTORY (Audit an LDAP directory server)

Use this command to audit an IBM Spectrum Protect™ controlled namespace on a Lightweight Directory Access Protocol (LDAP) server. The LDAP server and namespace are specified by using one or more LDAPURL options.

Restriction: Use this command only if you configured password authentication as described in Authenticating users by using an LDAP server. Information that is provided about the AUDIT LDAPDIRECTORY command applies only to environments in which

password authentication is configured as described in Authenticating users by using an LDAP server.

Nodes and administrator user IDs that do not authenticate their passwords with the LDAP directory server are deleted with the AUDIT LDAPDIRECTORY FIX=YES command. Nodes or administrator user IDs that no longer exist in the IBM Spectrum Protect database are also deleted.

Before you issue this command, ensure that the LDAPURL option is specified in the dsmserv.opt file. See the LDAPURL option for more information. If you specified more than one LDAPURL option in the dsmserv.opt file, each option is validated in the order in which they are placed. If the LDAPURL option is not specified, the command fails.

Privilege class

You must have system privileges to issue this command.

Syntax

```

>>-AUDIT LDAPdirectory--+-Fix-----No-----+----->
                          '-Fix-----+No--+-'
                          '-Yes-'

.-Wait-----No-----.
>--+-----+----->>
  '-Wait-----+No--+-'
  '-Yes-'
```

Parameters

Fix

This optional parameter specifies how the IBM Spectrum Protect server resolves inconsistencies between the database and the external directory. The default is NO. You can specify the following values:

No

The server reports all inconsistencies but does not change the external directory.

Yes

The server resolves any inconsistencies that it can and suggests further actions, if needed.

Important: If there are LDAP entries that are shared with other IBM Spectrum Protect servers, choosing YES might cause those servers to become out-of-sync.

Wait

This optional parameter specifies whether to wait for the IBM Spectrum Protect server to complete processing this command in the foreground. The default is NO. You can specify the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must complete before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Audit an LDAP directory and repair inconsistencies

Audit the LDAP directory that you specified in the LDAPURL option. The IBM Spectrum Protect server resolves some inconsistencies.

```
audit ldapdirectory fix=yes
```

```
ANR2749W Admin ADMIN1 was located in the LDAP directory server but not
in the database.
```

```
ANR2749W Admin ADMIN2 was located in the LDAP directory server but not
in the database.
```

ANR2749W Admin NODE1 was located in the LDAP directory server but not in the database.
 ANR2749W Admin NODE2 was located in the LDAP directory server but not in the database.
 ANR2748W Node NODE1 was located in the LDAP directory server but not in the database.
 ANR2748W Node NODE2 was located in the LDAP directory server but not in the database.
 ANR2745I AUDIT LDAPDIRECTORY command completed: 4 administrator entries are only in the LDAP directory server (not in the IBM Spectrum Protect server), 0 administrator entries are only in the IBM Spectrum Protect server (not in the LDAP directory server), 2 node entries are only in the LDAP directory server (not in the IBM Spectrum Protect server), 0 node entries are only in the IBM Spectrum Protect server, (not in the LDAP directory server), 6 entries were deleted from the LDAP server in total.

Related commands

Table 1. Commands related to AUDIT LDAPDIRECTORY

Command	Description
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.

AUDIT LIBRARY (Audit volume inventories in an automated library)

Use this command to audit and synchronize volume inventories in an automated library.

When the AUDIT LIBRARY command is issued on a library client, the client synchronizes its inventory with the inventory on the library manager. If the library client detects inconsistencies, it corrects them by changing the ownership of the volume on the library manager.

When the AUDIT LIBRARY command is issued on a server where the library is SCSI, 349X, or ACSLS (LIBTYPE=SCSI, LIBTYPE=349X, or LIBTYPE=ACSL), the server synchronizes its inventory with the inventory of the library device. If the server detects inconsistencies, it deletes missing volumes from its inventory.

- In SCSI libraries, the server also updates the locations of volumes in its inventory that have been moved since the last audit.
- In 349X libraries, the server also ensures that scratch volumes are in the scratch category and that private volumes are in the private category.

When the AUDIT LIBRARY command is issued on a server that is a library manager for the library (SHARED=YES), the server updates ownership of its volumes if it detects inconsistencies.

Regardless the type of server or type of library, issuing the AUDIT LIBRARY command does not automatically add new volumes to a library. To add new volumes, you must use the CHECKIN LIBVOLUME command.

Attention: The following precautions apply to SCSI, 349X, and ACSLS libraries only (LIBTYPE=SCSI, LIBTYPE=349X, and LIBTYPE=ACSL):

- Running the AUDIT LIBRARY command prevents any other library activity until the audit completes. For example, the server will not process restore or retrieve requests that involve the library when the AUDIT LIBRARY command is running.
- If other activity is occurring in the library, do not issue the AUDIT LIBRARY command. Issuing the AUDIT LIBRARY command when a library is active can produce unpredictable results (for example, a hang condition) if a process currently accessing the library attempts to acquire a new tape mount.

This command creates a background process that you can cancel with the CANCEL PROCESS command. To display information about background processes, use the QUERY PROCESS command.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-AUDIT LIBRARY--library_name----->
.-CHECKLabel----Yes-----
>--+-----+----->
'-CHECKLabel----+Yes----+'
      '-Barcode-'

.-REFRESHstate----No-----
>--+-----+----->>
'-REFRESHstate----+No----+'
      '-Yes-'
```

Parameters

library_name (Required)

Specifies the name of the library to audit.

CHECKLabel

Specifies how the storage volume label is checked during the audit. This parameter applies to SCSI libraries only. The parameter is ignored for other library types. The default is YES. Possible values are:

Yes

Specifies that the server checks each volume label to verify the identity of the volume.

Barcode

Specifies that the server uses the barcode reader to read the storage label. Using the barcode decreases the audit processing time. This parameter applies only to SCSI libraries.

Attention: If the scanner cannot read the barcode label or the barcode label is missing, the server loads that tape in a drive to read the label.

REFRESHstate

Specifies whether the server's information about a library, which is normally obtained during initialization, is refreshed, so that any changes in configuration are reflected. By setting the REFRESHSTATE parameter to Yes, this action is completed without having to restart the server or re-define the library. The default is No. Possible values are:

No

Specifies that the server does not refresh the library's state when the library is audited.

Yes

Specifies that the server does refresh the library's state when the AUDIT LIBRARY command is issued.

Example: Audit an automated library

Audit the EZLIFE automated library.

```
audit library ezlife
```

Related commands

Table 1. Commands related to AUDIT LIBRARY

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE LIBRARY	Deletes a library.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.

Command	Description
QUERY PROCESS	Displays information about background processes.
UPDATE LIBRARY	Changes the attributes of a library.

AUDIT LIBVOLUME (Verify database information for a tape volume)

Use this command to determine whether a tape volume is intact and to audit data on any tape volume.

You can issue the AUDIT LIBVOLUME command from any tape volume that is checked in to a library. The command runs in the background by default. You can issue the command from the following library types that have IBM® TS1140, IBM LTO 5, or a later generation tape drive:

- SCSI tape library
- Virtual tape library (VTL)

The following table outlines the tape drives that can verify tape volumes with media types for IBM TS1140 and IBM LTO 5 and later generation LTO tape drives:

Table 1. Tape drives and the media types

Drive	Media type
TS1140	JB, JX, JA, JW, JJ, JR, JC, JY, and JK
IBM LTO 5	LTO 3, LTO 4, and LTO 5
IBM LTO 6	LTO 4, LTO 5, and LTO 6
IBM LTO 7	LTO 5, LTO 6, and LTO 7

The following table outlines the minimum device driver level that you require to run the command:

Table 2. Minimum IBM device driver level

Driver name	Device driver level
Atape driver on AIX®	12.3.5.00
lin_tape driver on Linux	1.6.7.00
IBM tape driver on Windows	6.2.2.00

Restriction: You cannot issue the CANCEL PROCESS command while the AUDIT LIBVOLUME command is in progress.

Privilege class

To issue this command, you must have system privilege, or unrestricted storage privilege for the library to which the tape volume is defined.

Syntax

```
>>-AUDit LIBVolume--library_name--volume_name----->
      .-Wait-----No-----.
>--+-----+----->>
      '-Wait-----+No--+-'
              '-Yes-'
```

Parameters

library_name (Required)

Specifies the name of the library volume where the tape volume is located that you want to audit.

volume_name (Required)

Specifies the name of the physical tape volume that you want to audit.

Wait (Optional)

Specifies whether the audit or verification operation is completed in the foreground or background. This parameter is optional. The following options are available:

No

Specifies that the operation is completed in the background. The NO value is the default value.

Yes

Specifies that the operation is completed in the foreground. It might take a long time to complete the operation.

Example: Audit a tape volume

Audit the EZLIFE library that has a tape volume that is called KM0347L5.

```
audit libvolume ezlife KM0347L5
```

AUDIT LICENSES (Audit server storage usage)

Use this command to audit the server storage used by client nodes and to audit the server licenses. The audit determines whether the current configuration is in compliance with the license terms.

An audit creates a background process you can cancel with the CANCEL PROCESS command. If you halt and restart the server, an audit is run automatically as specified by the SET LICENSEAUDITPERIOD. To view audit results, use the QUERY LICENSE command.

Attention: The audit of server storage can take a lot of CPU time. You can use the AUDITSTORAGE server option to specify that storage is not to be audited.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-AUDit LICenses-----<<
```

Parameters

None.

Example: Audit server licenses

Issue the AUDIT LICENSES command.

```
audit licenses
```

Related commands

Table 1. Commands related to AUDIT LICENSES

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY AUDITOCCUPANCY	Displays the server storage utilization for a client node.
QUERY LICENSE	Displays information about licenses and audits.
QUERY PROCESS	Displays information about background processes.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the IBM Spectrum Protect server.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.

AUDIT VOLUME (Verify database information for a storage pool volume)

Use this command to check for inconsistencies between database information and a storage pool volume. Processing information generated during an audit is sent to the activity log and server console.

Restriction: You cannot use this command for volumes that are assigned to copy-container storage pools. You can only audit volumes that belong to storage pools with DATAFORMAT=NATIVE and DATAFORMAT=NONBLOCK.

You cannot audit a volume if it is being deleted from a primary or copy storage pool.

While an audit process is active, clients cannot restore data from the specified volume or store new data to that volume.

If the server detects a file with errors, handling of the file will depend on the type of storage pool to which the volume belongs, whether the FIX option is specified on this command, and whether the file is also stored on a volume assigned to other pools.

If IBM Spectrum Protect™ does not detect errors for a file that was marked as damaged, the state of the file is reset so that it can be used.

The server will not delete archive files that are on deletion hold. If archive retention protection is enabled, the server will not delete archive files whose retention period has not expired.

To display information about the contents of a storage pool volume, use the QUERY CONTENT command.

To audit multiple volumes, you can use the FROMDATE and TODATE parameters. Use the STGPOOL parameter to audit all volumes in a storage pool. When you use the parameters FROMDATE, TODATE, or both, the server limits the audit to only the sequential media volumes that meet the date criteria, and automatically includes all online disk volumes in storage. To limit the number of volumes that may include disk volumes, use the FROMDATE, TODATE, and STGPOOL parameters.

If you are running a server with archive retention protection enabled, and you have data stored in storage pools which are defined with the parameter RECLAMATIONTYPE=SNAPLOCK, the Last Access Date on the NetApp SnapLock Filer for a volume should be equal to the End Reclaim Period date that you see when you issue a QUERY VOLUME F=D command on that volume. During AUDIT VOLUME processing, these dates are compared. If they do not match and the AUDIT VOLUME command is being run with the FIX=NO parameter, a message will be issued to you indicating that the command should be run with the FIX=YES parameter to resolve the inconsistency. If they do not match and the AUDIT VOLUME command is being run with the FIX=YES parameter, the inconsistencies will be resolved.

Attention: Use the FIX=Yes parameter only if your tape drive and storage area network (SAN) infrastructure is stable. Ensure that the tape heads are clean and that the tape device drivers are stable and reliable. Otherwise, you risk deleting data that is error free when you use this parameter. The server cannot determine whether a tape is physically damaged or whether a tape infrastructure is unstable.

This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax

```

                                .-Fix-----No-----.
>>-AUDit Volume---+volume_name+-----+----->
                '-| A |-----' '-Fix-----+No---+'
                                '-Yes-'

    .-SKIPPartial-----No-----.  .-Quiet-----No-----.
>--+-----+-----+-----+-----><
    '-SKIPPartial-----+No---+' '-Quiet-----+No---+'
                                '-Yes-'          '-Yes-'

A (at least one of these parameters must be specified)

|-----+-----+-----+----->
| (1)                                     |
```



```

'-----STGPool---poolname-'
      (1)                                (1)
.-----FROMDate---TODAY-.  .-TODate---TODay-----
>--+-----+-----+-----+-----+-----+-----+-----+-----+-----+
'FROMDate---date-----'  'TODate---date-----'

```

Notes:

1. You cannot specify a volume name if you specify a storage pool name, FROMDATE, or TODATE.

Parameters

volume_name

Specifies the name of the storage pool volume you want to audit. This parameter is required if you do not specify a storage pool. You cannot specify a volume name together with the FROMDATE and TODATE parameters.

Fix

Specifies how the server resolves inconsistencies between the database inventory and the specified storage pool volume. This parameter is optional. The default is NO.

The actions the server performs depend on whether the volume is assigned to a primary or a copy storage pool.

Primary Storage Pool:

Note: If the AUDIT VOLUME command does not detect an error in a file that was previously marked as damaged, IBM Spectrum Protect resets the state of the file so that it can be used. This provides a means for resetting the state of damaged files if it is determined that the errors were caused by a correctable hardware problem such as a dirty tape head.

Fix=No

IBM Spectrum Protect reports, but does not delete, database records that refer to files with inconsistencies:

- IBM Spectrum Protect marks the file as damaged in the database. If a backup copy is stored in a copy storage pool, you can restore the file using the RESTORE VOLUME or RESTORE STGPOOL command.
- If the file is a cached copy, you must delete references to the file on this volume by issuing the AUDIT VOLUME command and specifying FIX=YES. If the physical file is not a cached copy, and a duplicate is stored in a copy storage pool, it can be restored by using the RESTORE VOLUME or RESTORE STGPOOL command.

Fix=Yes

The server fixes any inconsistencies as they are detected:

- If the physical file is a cached copy, the server deletes the database records that refer to the cached file. The primary file is stored on another volume.
- If the physical file is not a cached copy, and the file is also stored in one or more copy storage pools, the error will be reported and the physical file marked as damaged in the database. You can restore the physical file by using the RESTORE VOLUME or RESTORE STGPOOL command.
- If the physical file is not a cached copy, and the physical file is not stored in a copy storage pool, each logical file for which inconsistencies are detected are deleted from the database.
- If archive retention protection is enabled by using the SET ARCHIVERETENTIONPROTECTION command, a cached copy of data can be deleted if needed. Data in primary and copy storage pools can only be marked damaged and never deleted.

Do not use the AUDIT VOLUME command with FIX=YES if a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The AUDIT VOLUME command could cause the restore to be incomplete.

Copy Storage Pool:

Fix=No

The server reports the error and marks the physical file copy as damaged in the database.

Fix=Yes

The server deletes any references to the physical file and any database records that point to a physical file that does not exist.

SKIPPARTIAL

Specifies whether IBM Spectrum Protect ignores partial files, which are files that span multiple storage pool volumes. This parameter is optional. The default value is NO. When performing an audit operation on a sequential access media volume,

this parameter prevents additional sequential access media mounts that may be necessary to audit any partial files.
Possible values are:

No

IBM Spectrum Protect audits files that span multiple volumes.
Unless you specify SKIPPARTIAL=YES, IBM Spectrum Protect attempts to process each file stored on the volume, including files that span into and out of other volumes. To audit files that span multiple volumes, the following conditions must be true:

- For sequential access volumes, the additional sequential access volumes must have an access mode of read/write or read-only.
- For random access volumes, the additional volumes must be online.

Yes

IBM Spectrum Protect audits only files that are stored on the volume to be audited. The status of any partial files is unknown.

Quiet

Specifies whether IBM Spectrum Protect sends detailed informational messages to the activity log and the server console about irretrievable files on the volume. This parameter is optional. The default is NO. Possible values are:

No

Specifies that IBM Spectrum Protect sends detailed informational messages and a summary. Each message contains the node, file space, and client name for the file.

Yes

Specifies that IBM Spectrum Protect sends only a summary report.

FROMDate

Specifies the beginning date of the range to audit volumes. The default is the current date. All sequential media volumes meeting the time range criteria that were written to after this date are audited. The server includes all online disk volumes in storage. The server starts one audit process for each volume and runs the process serially. You cannot use this parameter if you have specified a volume. This parameter is optional. To limit the number of volumes that may include disk volumes, use the FROMDATE, TODATE, and STGPOOL parameters.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2001 If a date is entered, all candidate volumes written on that day (starting at 12:00:01 am) will be evaluated.
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 <i>or</i> -7. To display information beginning with volumes written a week ago, you can specify FROMDATE=TODAY-7 <i>or</i> FROMDATE= -7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

TODate

Specifies the ending date of the range for volumes to audit. All sequential media volumes meeting the time range criteria that were written to before this date are audited. The server includes all online disk volumes in storage. If you do not specify a value, the server defaults to the current date. You cannot use this parameter if you have specified a volume. This parameter is optional. To limit the number of volumes that may include disk volumes, use the FROMDATE, TODATE, and STGPPOOL parameters.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2001 If a date is entered, all candidate volumes written on that day (ending at 11:59:59 pm) will be evaluated.
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To display information created up to yesterday, you can specify TODATE=TODAY-1 or simply TODATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STGPool

This parameter specifies that the server only audits the volumes from the specified storage pool. This parameter is optional. You cannot use this parameter if you have specified a volume.

Example: Verify database information for a specific storage pool volume

Verify that the database information for storage pool volume PROG2 is consistent with the data stored on the volume. IBM Spectrum Protect fixes any inconsistencies.

```
audit volume prog2 fix=yes
```

Example: Verify database information for all volumes written to during a specific date range

Verify that the database information for all eligible volumes written to from 3/20/2002 to 3/22/2002 is consistent with data stored on the volume.

```
audit volume fromdate=03/20/2002 todate=03/22/2002
```

Example: Verify database information for all volumes in a specific storage pool

Verify that the database information for all volumes in storage pool STPOOL3 is consistent with data stored on the volume for today.

```
audit volume stgpool=STPOOL3
```

Example: Verify database information for all volumes in a specific storage pool written to in the last two days

Verify that the database information for all volumes in storage pool STPOOL3 is consistent with data stored on the volume for the last two days.

```
audit volume stgpool=STPOOL3 fromdate=-1
```

Related commands

Table 1. Commands related to AUDIT VOLUME

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
QUERY VOLUME	Displays information about storage pool volumes.
SET ARCHIVERETENTIONPROTECTION	Specifies whether data retention protection is activated.

BACKUP commands

Use the BACKUP commands to create backup copies of IBM Spectrum Protect™ information or objects.

- BACKUP DB (Back up the database)
- BACKUP DEVCONFIG (Create backup copies of device configuration information)
- BACKUP NODE (Back up a NAS node)
- BACKUP STGPOOL (Back up primary storage pool data to a copy storage pool)
- BACKUP VOLHISTORY (Save sequential volume history information)

BACKUP DB (Back up the database)

Use this command to back up an IBM Spectrum Protect™ database to sequential access volumes.

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

To determine how much extra storage space a backup requires, issue the QUERY DB command.

Restrictions: You cannot restore a server database if the release level of the server database backup is different from the release level of the server that is being restored. For example, an error occurs when you restore a Version 6.3 database and you are using a Version 7.1 server.

After the database backup is complete, the IBM Spectrum Protect server backs up information, depending on the options that are specified in the server options file. The following information is backed up:

- Sequential volume-history information is backed up to all files that the VOLUMEHISTORY option specifies
- Information about device configuration is backed up to all files that the DEVCONFIG option specifies
- The server's master encryption key

If there is not enough space available on the defined active log directory volume or file space, you can define the DB2® option, *overflowlogpath*, to use a directory with the required space available. For example, use the following command to use the `/home/tsminst2/overflow_dir` directory:

```
db2 update db cfg for TSMDb1 using overflowlogpath /home/tsminst2/overflow_dir
```

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-BACkup DB--DEVclass-----device_class_name----->
      .-Type-----Full-----+
>--+-----+-----+-----+-----+-----+----->
      '-Type-----+-Incremental-+-'
      +-Full-----+
```

```

'-DBSnapshot--'
>-----+-----+-----+-----+-----+----->
|
|          .,-----|.
|          V         |
|'-VOLumenames-----+---volume_name-----+---'
|          '-FILE:-- file_name-'
|
|. -NUMStreams-----1-----|.  |.-Scratch-----Yes-----|.
>-----+-----+-----+-----+-----+----->
|'-NUMStreams-----number-'  |'-Scratch-----+Yes--+-'
|                                |'-No--'
|
|. -Wait-----No-----|.  |.-DEDUPDEvice-----No-----|.
>-----+-----+-----+-----+-----+----->
|'-Wait-----+No--+-'  |'-DEDUPDEvice-----+No--+-'
|          '-Yes-'          |          '-Yes-'
|
|. -COMPRESS-----No-----|.  |.-PROTECTKeys-----Yes-----|.
>-----+-----+-----+-----+-----+----->
|          (1) |  |'-PROTECTKeys-----+No--+-'
|'-COMPRESS-----+No--+-'          |'-Yes-'
|          '-Yes-'
|
>-----+-----+-----+-----+-----+----->>
|'-PASSword-----password_name-'

```

Notes:

1. The default value of the COMPRESS parameter is conditional. If you specify the COMPRESS parameter in the BACKUP DB command, it overrides any COMPRESS parameter value that is set in the SET DBRECOVERY command. Otherwise, the value that is set in the SET DBRECOVERY command is the default.

Parameters

DEVclass (Required)

Specifies the name of the sequential access device class to use for the backup. If you issue the BACKUP DB command, and the device class is not the one that is specified in the SET DBRECOVERY command, a warning message is issued. However, the backup operation continues and is not affected.

If the SET DBRECOVERY command is not issued to set a device class, the BACKUP DB command fails.

Restriction:

- You cannot use a device class with a device type of NAS or CENTERA.
- A restore database operation fails if the source for the restore is a FILE library. A FILE library is created if the FILE device class specifies SHARED=YES.

If all drives for this device class are busy when the backup runs, IBM Spectrum Protect cancels lower priority operations, such as reclamation, to make a drive available for the backup.

Type

Specifies the type of backup to run. This parameter is optional. The default is FULL. The following values are possible:

Full

Specifies that you want to run a full backup of the IBM Spectrum Protect database.

Incremental

Specifies that you want to run an incremental backup of the IBM Spectrum Protect database. An incremental (or cumulative) backup image contains a copy of all database data that is changed since the last successful full backup operation.

DBSnapshot

Specifies that you want to run a full snapshot database backup. The entire contents of a database are copied and a new snapshot database backup is created without interrupting the existing full and incremental backup series for the database.

VOLumenames

Specifies the volumes that are used to back up the database. This parameter is optional. However, if you specify SCRATCH=NO, you must specify a list of volumes.

volume_name

Specifies the volumes that are used to back up the database. Specify multiple volumes by separating the names with commas and no intervening spaces.

FILE:filename

Specifies the name of a file that contains a list of volumes that are used to back up the database. Each volume name must be on a separate line. Blank lines and comment lines, which begin with an asterisk, are ignored.

For example, to use volumes DB0001, DB0002, and DB0003, create a file that contains these lines:

```
DB0001
DB0002
DB0003
```

Name the file appropriately. For example:

- **AIX** | **Linux** TAPEVOL
- **Windows** TAPEVOL.DATA

You can then specify the volumes for the command as follows:

AIX | **Linux**

```
VOLUMENAMES=FILE:TAPEVOL
```

Windows

```
VOLUMENAMES=FILE:TAPEVOL.DATA
```

NUMStreams

Specifies the number of parallel data movement streams to use when you back up the database. The minimum value is 1, and the maximum value is 32. Increasing the value causes a corresponding increase in the number of database backup sessions to be used and the number of drives to be used for the device class. If you specify a NUMSTREAMS value in the BACKUP DB command, it overrides any value that is set in the SET DBRECOVERY command. Otherwise, the value that is set in the SET DBRECOVERY command is used. The NUMSTREAMS value is used for all types of database backups.

If a value is specified that is greater than the number of drives available for the device class, only the number of available drives are used. The available drives are those defined to the device class by the MOUNTLIMIT parameter or by the number of online drives for the specified device class. The session is displayed in the QUERY SESSION output.

If you increase the number of streams, more volumes are used from the corresponding device class for this operation.

Using more volumes might improve the speed of the database backups, but at the cost of more volumes that are not fully used.

Scratch

Specifies whether scratch volumes can be used for the backup. This parameter is optional. The default is YES. The following values are possible:

Yes

Specifies that scratch volumes can be used.

If you specify SCRATCH=YES and the VOLUMENAMES parameter, IBM Spectrum Protect uses only scratch volumes if space is unavailable on the specified volumes.

If you do not include a list of volumes by using the VOLUMENAMES parameter, you must either specify SCRATCH=YES or use the default.

No

Specifies that scratch volumes cannot be used.

If you specify volumes by using the VOLUMENAMES parameter and SCRATCH=NO, the backup fails if there is not enough space available to store the backup data on the specified volumes.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is NO. The following values are possible:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If a BACKUP DB background process is canceled, some of the database might have already been backed up before the cancellation.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before you continue with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

DEDUPDEvice

Specifies that a target storage device supports data deduplication. When set to YES, the format for backup images is optimized for data deduplication devices, making backup operations more efficient. The following values are possible:

No

Specifies that a target storage device does not support data deduplication. NO is the default.

Ensure that this parameter is set to NO for the following devices:

- SCSI libraries
- All devices that are defined with a FILE device class
- Virtual tape libraries (VTL) that do not support the data deduplication function

Yes

Specifies that a target device supports data deduplication and that you want to optimize backups for this function. You can set this parameter to YES if you are using VTLs that support data deduplication.

COMPRESS

Specifies whether volumes that are created by the BACKUP DB command are compressed. The COMPRESS value is used for all types of database backups. This parameter is optional. The default value is conditional. If you specify the COMPRESS parameter on the BACKUP DB command, it overrides any value that is set in the SET DBRECOVERY command. Otherwise, the value that is set in the SET DBRECOVERY command is the default. You can specify one of the following values:

No

Specifies that the volumes that are created by the BACKUP DB command are not compressed.

Yes

Specifies that the volumes that are created by the BACKUP DB command are compressed.

Restrictions:

- Use caution when you specify the COMPRESS parameter. Using compression during database backups can reduce the size of the backup files. However, compression can increase the time that is required to complete database backup processing.
- Do not back up compressed data to tape. If your system environment stores database backups on tape, set the COMPRESS parameter to No in the SET DBRECOVERY and BACKUP DB commands.

AIX	Linux	Windows	PROTECTKeys
-----	-------	---------	-------------

AIX	Linux	Windows	Specifies that database backups include a copy of the server master encryption key that is used to encrypt node passwords, administrator passwords, and storage pool data. The master encryption key is stored in the dsmkeydb files. If you lose the dsmkeydb files, nodes and administrators are unable to authenticate with the server because the server is unable to read the passwords that are encrypted by using the master encryption key. In addition, any data that is stored in an encrypted storage pool cannot be retrieved without the master encryption key. This parameter is optional. The default is the value that is specified for the PROTECTKEYS parameter on the SET DBRECOVERY command. You can specify one of the following values:
-----	-------	---------	---

No

Specifies that database backups do not include a copy of the server master encryption key.

Attention: If you specify PROTECTKEYS=NO, you must manually back up the master encryption key for the server and make the key available when you implement disaster recovery. You cannot recover from a disaster without the master encryption key.

Yes

Specifies that database backups include a copy of the server master encryption key.

Attention: If you specify PROTECTKEYS=YES, you must also specify the PASSWORD parameter.

AIX | **Linux** | **Windows** | **PASS**word

AIX | **Linux** | **Windows** Specifies the password that is used to protect the database backup. The default is the value that is specified for the PASSWORD parameter on the SET DBRECOVERY command. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

Important: Ensure that you remember this password. If you specify a password for database backups, you must specify the same password on the RESTORE DB command to restore the database.

Example: Run an incremental backup by using a scratch volume

Run an incremental backup of the database by using a scratch volume. Use a device class of FILE for the backup.

```
backup db devclass=file type=incremental
```

AIX | **Linux** | **Windows**

Example: Encrypt storage pool data in database backups

Encrypt storage pool data by specifying that database backups include a copy of the server master encryption key. Issue the following command:

```
backup db protectkeys=yes password=password_name
```

Related commands

Table 1. Commands related to BACKUP DB

Command	Description
BACKUP DEVCONFIG	Backs up IBM Spectrum Protect device information to a file.
BACKUP VOLHISTORY	Records volume history information in external files.
CANCEL PROCESS	Cancels a background server process.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
PREPARE	Creates a recovery plan file.
QUERY DB	Displays allocation information about the database.
QUERY PROCESS	Displays information about background processes.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DBRECOVERY	Specifies the device class to be used for automatic backups.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.

BACKUP DEVCONFIG (Create backup copies of device configuration information)

Use this command to back up information about device configuration for the server.

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

This command backs up the following information in one or more files:

- Device class definitions
- Library definitions

- Drive definitions
- Path definitions when SRCTYPE=SERVER
- Server definitions
- Server name
- Server password
- Volume location information for LIBTYPE=SCSI libraries

AIX | **Linux** You can use the DEVCONFIG server option to specify one or more files in which to store device configuration information. IBM Spectrum Protect™ updates the files whenever a device class, library, or drive is defined, updated, or deleted.

Windows At installation, the server options file includes a DEVCONFIG option that specifies a device configuration file named devcnfg.out. IBM Spectrum Protect updates this file whenever a device class, library, or drive is defined, updated, or deleted.

To ensure updates are complete before the server is halted:

- Do not halt the server for a few minutes after issuing the BACKUP DEVCONFIG command.
- Specify multiple DEVCONFIG options in the server options file.
- Examine the device configuration file to see if the file has been updated.

Privilege class

Any administrator can issue this command unless it includes the FILENAMES parameter. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES, the administrator must have system privilege. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, policy, storage or system privilege.

Syntax

```
>>-Backup DEVCONFig-----+-----+----->>
|                               .-+-----+ |
|                               V          | |
|'-Filenames-----filename-----+'
```

Parameters

Filenames

Specifies the files in which to store device configuration information. You can specify multiple files by separating the names with commas and no intervening spaces. This parameter is optional.

If you do not specify a file name, IBM Spectrum Protect stores the information in all files specified with the DEVCONFIG option in the server options file.

Example: Backup device configuration information to a file

Back up device configuration information to a file named DEVICE.

```
backup devconfig filenames=device
```

Related commands

Table 1. Commands related to BACKUP DEVCONFIG

Command	Description
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DEVCLASS	Defines a device class.
AIX Linux DEFINE DEVCLASS (z/OS® media server)	AIX Linux Defines a device class to use storage managed by a z/OS media server.
DEFINE DRIVE	Assigns a drive to a library.

Command	Description
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBVOLUME	Displays information about a library volume.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERPASSWORD	Specifies the server password.
UPDATE DEVCLASS	Changes the attributes of a device class.
AIX Linux UPDATE DEVCLASS (z/OS media server)	AIX Linux Changes the attributes of a device class for storage managed by a z/OS media server.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE PATH	Changes the attributes associated with a path.
UPDATE SERVER	Updates information about a server.

BACKUP NODE (Back up a NAS node)

Use this command to start a backup operation for a network-attached storage (NAS) node.

Backups that are created for NAS nodes with this BACKUP NODE command are functionally equivalent to backups that are created by using the BACKUP NAS command on an IBM Spectrum Protect™ client. You can restore these backups with either the server's RESTORE NODE command or the client's RESTORE NAS command.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax

```
>>-BBackup Node--node_name--+-----+----->
      | .-,------. |
      | v              | |
      '-----file_system_name-----'

      .-TOC-----Preferred-----.
>--+-----+-----+----->
  '-Mgmtclass-----mcname-' '-TOC-----+No-----+'
                                   +-Preferred+
                                   '-Yes-----'

  .-Wait-----No-----.  .-MODE-----DIFFerential-----.
>--+-----+-----+----->
  '-Wait-----+No--+-' '-MODE-----+FULL-----+-'
      '-Yes-'                '-DIFFerential-'

  .-TYPE-----BACKUPImage-----.
>--+-----+-----+----->>
  '-TYPE-----+BACKUPImage+-'
      '-SNAPMirror--'
```

Parameters

node_name (Required)

Specifies the node for which the backup will be performed. You cannot use wildcard characters or specify a list of names.

file_system_name

Specifies the name of one or more file systems to back up. You can also specify names of virtual file spaces that have been defined for the NAS node. The file system name that you specify cannot contain wildcard characters. You can specify more than one file system by separating the names with commas and no intervening spaces.

If you do not specify a file system, all file systems will be backed up. Any virtual file spaces defined for the NAS node are backed up as part of the file system image, not separately.

If a file system exists on the NAS device with the same name as the virtual file space specified, IBM Spectrum Protect automatically renames the existing virtual file space in the server database, and backs up the NAS file system which matches the name specified. If the virtual file space has backup data, the file space definition associated with the virtual file space will also be renamed.

Tip: See the virtual file space name parameter in the DEFINE VIRTUALFSMAPPING command for more naming considerations.

In determining the file systems to process, the server will not use any DOMAIN.NAS, INCLUDE.FS.NAS, or EXCLUDE.FS.NAS statements in any client option file or client option set. If you back up multiple file systems, the backup of each file system is a separate server process.

MGmtclass

Specifies the name of the management class to which this backup data is bound. If you do not specify a management class, the backup data is bound to the default management class of the policy domain to which the node is assigned. In determining the management class, the server will *not* use any INCLUDE.FS.NAS statements in any client option file or client option set. The destination management class might refer to an IBM Spectrum Protect native pool, in which case Network Data Management Protocol (NDMP) data is sent into the IBM Spectrum Protect native hierarchy. After this occurs, the data stays in the IBM Spectrum Protect hierarchy. Data flowing to IBM Spectrum Protect native pools goes over the LAN and data flowing to NAS pools can be directly attached or over a SAN.

When you specify a management class with the BACKUP NODE command, all versions of the backup data that belong to the NAS node are rebound to the new management class.

TOC

Specifies whether a table of contents (TOC) is saved for each file system backup. Consider the following in determining whether you want to save a table of contents:

- If a table of contents is saved, you will be able to use the QUERY TOC command to determine the contents of a file system backup in conjunction with the RESTORE NODE command to restore individual files or directory trees. You can also use the IBM Spectrum Protect web backup-archive client to examine the entire file system tree and choose files and directories to restore. Creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Creating a table of contents requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.
- A table of contents for a NAS file system cannot have a directory path greater than 1024 characters.
- If a table of contents is not saved for a file system backup, you will still be able to restore individual files or directory trees using the RESTORE NODE command, provided that you know the fully qualified name of each file or directory to be restored and the image in which that object was backed up.

This parameter is optional. The default value is Preferred. Possible values are:

No

Specifies that table of contents information is not saved for file system backups.

Preferred

Specifies that table of contents information should be saved for file system backups. However, a backup does not fail just because an error occurs during creation of the table of contents. This is the default value.

Yes

Specifies that table of contents information must be saved for each file system backup. A backup fails if an error occurs during creation of the table of contents.

Attention: If MODE=DIFFERENTIAL is specified and a table of contents is requested (TOC=PREFERRED or TOC=YES), but the last full image does not have a table of contents, a full backup will be performed and a table of contents will be created for that full backup.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is NO.
Possible values are:

No

Specifies that the server processes this command in the background. Use the QUERY PROCESS command to monitor the background processing of this command.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes. If you are backing up multiple file systems, all backup processes must complete before the command is complete.

Attention: You cannot specify WAIT=YES from the server console.

MODE

Specifies whether the file system backups are full or differential. The default is DIFFERENTIAL.

FULL

Specifies to back up the entire file system.

DIFFerential

Specifies that only the files that have changed since the most recent full backup should be backed up. If you choose a differential backup, and a full backup is not found, a full backup is performed. You cannot specify TYPE=SNAPMIRROR when the MODE parameter is set to DIFFERENTIAL.

TYPE

Specifies the backup method used to perform the NDMP backup operation. The default value for this parameter is BACKUPIMAGE and it should be used to perform a standard NDMP base or differential backup. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

BACKUPImage

Specifies that the file system should be backed up using an NDMP dump operation. This is the default method for performing an NDMP backup. The BACKUPIMAGE type operation supports full and differential backups, file-level restore processing and directory-level backup.

SNAPMirror

Specifies that the file system should be copied to an IBM Spectrum Protect storage pool using the NetApp SnapMirror to Tape function. SnapMirror images are block level full backup images of a file system. Typically, a SnapMirror backup takes significantly less time to perform than a traditional NDMP full file system backup. However there are limitations and restrictions on how SnapMirror images can be used. The SnapMirror to Tape function is intended to be used as a disaster-recovery option for copying very large NetApp file systems to secondary storage.

For most NetApp file systems, use the standard NDMP full or differential backup method. Refer to the documentation that came with your NetApp file server for more information.

When setting the TYPE parameter to SNAPMirror, the following restrictions apply:

Restrictions:

- You cannot specify TOC=YES or TOC=PREFERRED.
- The file_system_name cannot be a virtual file space name.
- The snapshot which is created automatically by the file server during the SnapMirror copy operation will be deleted at end of the operation.
- This parameter is valid for NetApp and IBM® N-Series file servers only.

Example: Perform a full backup

Perform a full backup on the /vol/vol10 file system of NAS node NAS1.

```
backup node nas1 /vol/vol10 mode=full
```

Example: Perform a backup on a directory and create a table of contents

Back up the directory /vol/vol2/mikes on the node NAS1 and create a table of contents for the image. For the following two examples, assume Table 1 contains the virtual file space definitions exist on the server for the node NAS1.

```
backup node nas1 /mikesdir
```

Table 1. Virtual file space definitions

Virtual file space name	File system	Path
/mikesdir	/vol/vol2	/mikes
/DataDirVol2	/vol/vol2	/project1/data
/TestDirVol1	/vol/vol1	/project1/test

Example: Perform a backup on two directories

Back up the directories /vol/vol2/project1/data and /vol/vol1/project1/test of the node NAS1. Refer to Table 1 for the virtual file space definitions that exist on the server for the node NAS1.

```
backup node nas1 /DataDirVol2,/testdirvol1 mode=full toc=yes
```

Related commands

Table 2. Commands related to BACKUP NODE

Command	Description
BACKUP NAS (client command)	Creates a backup of NAS node data.
CANCEL PROCESS	Cancels a background server process.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY NASBACKUP	Displays information about NAS backup images.
QUERY TOC	Displays details about the table of contents for a specified backup image.
QUERY COPYGROUP	Displays the attributes of a copy group.
RESTORE NAS (client command)	Restores a backup of NAS node data.
RESTORE NODE	Restores a network-attached storage (NAS) node.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

Related concepts:

Backup and restore using NetApp SnapMirror to Tape feature

BACKUP STGPOOL (Back up primary storage pool data to a copy storage pool)

Use this command to back up primary storage pool files to a copy storage pool.

You can back up data from a primary storage pool that is defined with the NATIVE, NONBLOCK, or any of the NDMP formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The copy storage pool to which data is to be backed up must have the same data format as the primary storage pool. IBM Spectrum Protect™ supports back-end data movement for NDMP images.

If a file exists in the copy storage pool, the file is not backed up unless the copy of the file in the copy storage pool is marked as damaged. However, a new copy is not created if the file in the primary storage pool is also marked as damaged. In a random-access storage pool, cached copies of migrated files and damaged primary files are not backed up.

Tip: Issuing this command for a primary storage pool that is set up for data deduplication removes duplicate data, if the copy storage pool is also set up for data deduplication.

If migration for a storage pool starts during a storage pool backup, some files might be migrated before they are backed up. You might want to back up storage pools that are higher in the migration hierarchy before you back up storage pools that are lower.

Restrictions:

- Do not run the MOVE DRMEDIA and BACKUP STGPOOL commands concurrently. Ensure that the storage pool backup processes are complete before you issue the MOVE DRMEDIA command.
- You cannot back up data from or to storage pools defined with a CENTERA device class.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the copy storage pool in which backup copies are to be produced.

Syntax

```
>>-BBackup STGpool--primary_pool_name--copy_pool_name----->
. -MAXPRocess-----1----- .
>--+-----+----->
' -MAXPRocess-----number-'

. -Preview-----No----- .
>--+-----+----->
' -Preview-----+No-----+'
      +-Yes-----+
      |                (1) |
      '-VOLumesonly-----'

. -SHREDTONOshred-----No----- .   . -Wait-----No----- .
>--+-----+----->>
' -SHREDTONOshred-----+No--+-'   ' -Wait-----+No--+-'
      '-Yes-'                       '-Yes-'
```

Notes:

1. Valid only for storage pools that are associated with a sequential-access device class.

Parameters

primary_pool (Required)

Specifies the primary storage pool.

copy_pool (Required)

Specifies the copy storage pool.

MAXPRocess

Specifies the maximum number of parallel processes to use for backing up files. This parameter is optional. Enter a value 1 - 999. The default is 1.

Using multiple, parallel processes can improve throughput for the backup. The expectation is that the time needed to complete the storage pool backup is decreased by using multiple processes. However, when multiple processes are running, in some cases one or more of the processes needs to wait to use a volume that is already in use by a different backup process.

When you determine this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the backup.

Each process needs a mount point for copy storage pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are backing up a sequential storage pool, each process needs an extra mount point for primary storage pool volumes and, if the device type is not FILE, an extra drive. For example, suppose that you specify a maximum of three processes to back up a primary sequential storage pool to a copy storage pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least 6, and at least six mount points and six drives must be available.

To preview a backup, only one process is used and no mount points or drives are needed.

Preview

Specifies whether you want to preview but not run the backup. The preview displays the number of files and bytes to be backed up and a list of the primary storage pool volumes that you must mount. This parameter is optional. The default is NO. You can specify the following values:

No

Specifies that the backup is done.

Yes

Specifies that you want to preview the backup but not do the backup.

VOLUMESonly

Specifies that you want to preview the backup only as a list of the volumes that must be mounted. This choice requires the least processing time. The VOLUMESONLY option is valid only for storage pools that are associated with a sequential-access device class.

The VOLUMESONLY option can be used to obtain a list of volumes that are needed by the storage pool backup process. For example:

```
backup stgpool primary_pool copystg preview=volumesonly
```

The list of volumes are logged in the server activity log with the ANR1228I message. Query the server activity log to get the list of volumes required. For example:

```
query actlog msg=1228
```

SHREDTONOshred

Specifies whether data is backed up to a copy storage pool from a primary storage pool that enforces shredding. This parameter is optional. The default value is NO. You can specify the following values:

No

Specifies that the server does not allow data to be backed up to a copy storage pool from a primary storage pool that enforces shredding. If the primary storage pool enforces shredding, the operation fails.

Yes

Specifies that the server does allow data to be backed up to a copy storage pool from a primary storage pool that enforces shredding. The data in the copy storage pool is not shredded when it is deleted.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. You can specify the following values:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files might already have been backed up before the cancellation.

Yes

Specifies that the server processes this operation in the foreground. You must wait for the operation to complete before you continue with other tasks. The server displays the output messages to the administrative client when the operation completes.

Note: You cannot specify WAIT=YES from the server console.

Example: Back up the primary storage pool

Back up the primary storage pool that is named PRIMARY_POOL to the copy storage pool named COPYSTG.

```
backup stgpool primary_pool copystg
```

Related commands

Table 1. Commands related to BACKUP STGPOOL

Command	Description
CANCEL PROCESS	Cancel a background server process.

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
SHRED DATA	Manually starts the process of shredding deleted data.

BACKUP VOLHISTORY (Save sequential volume history information)

Use this command to back up sequential volume history information to one or more files.

Tip: You must use volume history information when you reload the database and audit affected storage pool volumes. If you cannot start the server, you can use the volume history file to query the database about these volumes.

The volume history includes information about the following types of volumes:

- Archive log volumes
- Database backup volumes
- Export volumes
- Backup set volumes
- Database snapshot volumes
- Database recovery plan file volumes
- Recovery plan file volumes
- Recovery plan file snapshot volumes
- The following sequential access storage pool volumes:
 - Volumes added to storage pools
 - Volumes reused through reclamation or MOVE DATA operations
 - Volumes removed by using the DELETE VOLUME command or during reclamation of scratch volumes

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

AIX | **Linux** You must use the VOLUMEHISTORY server option to specify one or more volume history files. IBM Spectrum Protect™ updates volume history files whenever server sequential volume history information is changed.

Windows At installation, the server options file includes a VOLUMEHISTORY option that specifies a default volume history file named volhist.out. IBM Spectrum Protect updates volume history files whenever server sequential volume history information is changed.

To ensure that updates are complete before the server is halted, follow these steps:

- Do not halt the server for a few minutes after you issue the BACKUP VOLHISTORY command.
- Specify multiple VOLUMEHISTORY options in the server options file.
- Examine the volume history file to see if the file has been updated.

Privilege class

Any administrator can issue this command unless it includes the FILENAMES parameter. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES, the administrator must have system privilege. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, policy, storage or system privilege.

Syntax


```
>>-Backup VOLHistory----->>
| .,-----|
| V          |
|-----file_name-----|
'-Filenames-----'
```

Parameters

Filenames

Specifies the names of one or more files in which to store a backup copy of volume history information. Separate multiple file names with commas and no intervening spaces. This parameter is optional.

If you do not specify a file name, IBM Spectrum Protect stores the information in all files specified with the VOLUMEHISTORY option in the server options file.

Example: Back up the volume history information to a file

Back up the volume history information to a file called VOLHIST.

```
backup volhistory filenames=volhist
```

Related commands

Table 1. Commands related to BACKUP VOLHISTORY

Command	Description
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
DELETE VOLUME	Deletes a volume from a storage pool.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.

BEGIN EVENTLOGGING (Begin logging events)

Use this command to begin logging events to one or more receivers. A receiver for which event logging has begun is an *active receiver*.

When the server is started, event logging automatically begins for the console and activity log and for any receivers that are started automatically based on entries in the server options file. You can use this command to begin logging events to receivers for which event logging is *not* automatically started at server startup. You can also use this command after you have disabled event logging to one or more receivers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-BEGin EVentlogging----->>
| .,-----|
| V          |
|-----|
|-----CONSOLE-----|
|-----+
|+--ACTLOG-----+
|+--EVENTSERVER----+
|+--FILE-----+
|+--FILETEXT-----+
|          (1) |
|+--NTEVENTLOG-----+
|-----|
```

```

|           (2)           |
+--SYSLOG-----+
+--TIVOLI-----+
'-USEREXIT-----'

```

Notes:

1. This parameter is only available for the Windows operating system.
2. This parameter is only available for the Linux operating system.

Parameters

Specify one or more receivers. You can specify multiple receivers by separating them with commas and no intervening spaces. If you specify ALL, logging begins for all receivers that are configured. The default is ALL.

ALL

Specifies all receivers that are configured for event logging.

CONSOLE

Specifies the server console as a receiver.

ACTLOG

Specifies the IBM Spectrum Protect™ activity log as a receiver.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

Windows NTEVENTLOG

Specifies the Windows application log as a receiver.

Linux SYSLOG

Specifies the Linux system log as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which IBM Spectrum Protect writes information as a receiver.

Example: Begin logging events

Begin logging events to the IBM Spectrum Protect activity log.

```
begin eventlogging actlog
```

Related commands

Table 1. Commands related to BEGIN EVENTLOGGING

Command	Description
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

CANCEL commands

Use the CANCEL commands to end a task or process before it is completed.

- CANCEL EXPIRATION (Cancel an expiration process)
- CANCEL EXPORT (Delete a suspended export operation)
- CANCEL PROCESS (Cancel an administrative process)
- CANCEL REPLICATION (Cancel node replication processes)
- CANCEL REQUEST (Cancel one or more mount requests)
- CANCEL RESTORE (Cancel a restartable restore session)
- CANCEL SESSION (Cancel one or more client sessions)

CANCEL EXPIRATION (Cancel an expiration process)

Use this command to cancel a process with an unknown process number that is running as a result of an inventory expiration operation.

Use the CANCEL EXPIRATION command if the expiration process number is not known, otherwise use the CANCEL PROCESS and specify the process number of the expiration process. Both commands call the same code to end the expiration process.

You can use the CANCEL EXPIRATION command to automate the cancellation of an expiration process. For example, if you start inventory expiration at midnight and, due to the maintenance workload on the server, the process must finish at 03:00, you can schedule a CANCEL EXPIRATION command to run at 03:00 without knowing the process number.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-CANcel EXPIration-----><
```

Example: Cancel an inventory expiration process

Cancel the process that was generated by an inventory expiration operation.

```
cancel expiration
```

Related commands

Table 1. Command related to CANCEL EXPIRATION

Command	Description
QUERY PROCESS	Displays information about background processes.
EXPIRE INVENTORY	Manually starts inventory expiration processing.

CANCEL EXPORT (Delete a suspended export operation)

Use this command to delete a suspended server-to server export operation. After issuing the CANCEL EXPORT command, you cannot restart the export operation. Issue the CANCEL PROCESS command to delete a currently running export operation.

Privilege class

You must have system privilege to issue this command.

Syntax

```
>>-CANcel EXPort .-*-----+-----><
                  +-----+-----><
                  '---export_identifier---'
```

Parameters

export_identifier

The unique identifier of the suspended export operation that you wish to delete. You can also enter wildcard characters for the identifier. Issue the QUERY EXPORT command to list the currently suspended export operations.

Example: Delete a specific suspended export operation

Cancel the suspended server-to-server export operation EXPORTALLACCTNODES.

```
cancel export exportallacctnodes
```

Example: Delete all suspended server-to-server export operations

Cancel all suspended server-to-server export processes.

```
cancel export *
```

Related commands

Table 1. Commands related to CANCEL EXPORT

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

CANCEL PROCESS (Cancel an administrative process)

Use this command to cancel a background process started by an administrative command or by a process, such as storage pool migration.

The following commands generate background processes:

- AUDIT CONTAINER
- AUDIT LIBRARY
- AUDIT LICENSES
- AUDIT VOLUME
- BACKUP DB
- BACKUP NODE
- BACKUP STGPOOL
- CHECKIN LIBVOLUME
- CHECKOUT LIBVOLUME
- **AIX** **Linux** **Windows** CONVERT STGPOOL
- DELETE FILESPACE
- DELETE VOLUME
- EXPIRE INVENTORY
- EXPORT ADMIN
- EXPORT NODE
- EXPORT POLICY
- EXPORT SERVER
- GENERATE BACKUPSET
- IMPORT ADMIN

- IMPORT NODE
- IMPORT POLICY
- IMPORT SERVER
- MIGRATE STGPOOL
- MOVE DATA
- MOVE DRMEDIA
- MOVE MEDIA
- PREPARE
- PROTECT STGPOOL
- RECLAIM STGPOOL
- REPLICATE NODE
- RESTORE NODE
- RESTORE STGPOOL
- RESTORE VOLUME
- VARY

The following internal server operations generate background processes:

- Inventory expiration
- Migration
- Reclamation

To cancel a process, you must have the process number, which you can obtain by issuing the QUERY PROCESS command.

Some processes, such as reclamation, generate mount requests to complete processing. If a process has a pending mount request, the process might not respond to a CANCEL PROCESS command until the mount request is answered or canceled by using the REPLY or CANCEL REQUEST command, or by timing out.

Issue the QUERY REQUEST command to list open requests, or query the activity log to determine whether a process has a pending mount request. A mount request indicates that a volume is needed for the current process, but the volume is not available in the library. The volume might not be available if the administrator issues the MOVE MEDIA or CHECKOUT LIBVOLUME command, or manually removes the volume from the library.

After you issue a CANCEL PROCESS command for an export operation, the process cannot be restarted. To stop a server-to-server export operation but allow it to be restarted later, issue the SUSPEND EXPORT command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-CANcel PRocess--process_number-----<<
```

Parameters

process_number (Required)
Specifies the number of the background process you want to cancel.

Example: Cancel a background process by using its process number

Cancel background process number 3.

```
cancel process 3
```

Related commands

Table 1. Commands related to CANCEL PROCESS

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.

Command	Description
CANCEL REQUEST	Cancels pending volume mount requests.
AIX Linux Windows CONVERT STGPOOL	AIX Linux Windows Convert a storage pool to a directory-container storage pool.
AIX Linux Windows PROTECT STGPOOL	AIX Linux Windows Protects a directory-container storage pool.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
REPLY	Allows a request to continue processing.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

CANCEL REPLICATION (Cancel node replication processes)

Use this command to cancel all node replication processes.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-CANcel REPLication-----<<
```

Parameters

None.

Example: Cancel node replication processes

Cancel all node replication processes.

```
cancel replication
```

Related commands

Table 1. Commands related to CANCEL REPLICATION

Command	Description
QUERY PROCESS	Displays information about background processes.
QUERY REPLICATION	Displays information about node replication processes.

CANCEL REQUEST (Cancel one or more mount requests)

Use this command to cancel one or more pending media mount requests. To cancel a mount request, you need to know the request number assigned to the request. This number is included in the mount request message and can also be shown by using the QUERY REQUEST command.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
>>-CANcel REQuest--+-request_number-+-----><
                '-All-----' '-PERManent-'
```

Parameters

request_number

Specifies the request number of the mount request to cancel.

ALL

Specifies to cancel all pending mount requests.

PERManent

Specifies that you want the server to flag the volumes for which you are canceling a mount request as unavailable. This parameter is optional.

Example: Cancel a mount request

Cancel request number 2.

```
cancel request 2
```

Related commands

Table 1. Commands related to CANCEL REQUEST

Command	Description
QUERY REQUEST	Displays information about all pending mount requests.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

CANCEL RESTORE (Cancel a restartable restore session)

Use this command to cancel a restartable restore session. You can cancel restore sessions in the active or restartable state. Any outstanding mount requests related to this session are automatically canceled.

To display restartable restore sessions, use the QUERY RESTORE command.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

```
>>-CANcel--REStore--+-session_number-+-----><
                '-All-----'
```

Parameters

session_number

Specifies the number for the restartable restore session. An active session is a positive number, and a restartable session is a negative number.

ALL

Specifies that all the restartable restore sessions are to be canceled.

Example: Cancel restore operations

Cancel all restore operations.

```
cancel restore all
```

Related commands

Table 1. Commands related to CANCEL RESTORE

Command	Description
QUERY RESTORE	Displays information about restartable restore sessions.

CANCEL SESSION (Cancel one or more client sessions)

Use this command to cancel existing administrative or client node sessions, and to force an administrative or client node session off the server. Any outstanding mount requests related to this session are automatically canceled. The client node must start a new session to resume activities.

If you cancel a session that is in the idle wait (IdleW) state, the client session is automatically reconnected to the server when it starts to send data again.

If this command interrupts a process, such as backup or archive, the results of any processing active at the time of interruption are rolled back and not committed to the database.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

```
>>-CANcel SEssion--+-session_number-+-----><
                    '-ALL-----'
```

Parameters

session_number

Specifies the number of the administrative, server, or client node sessions that you want to cancel.

ALL

Specifies that all client node sessions are canceled. You cannot use this parameter to cancel administrative client or server sessions.

Example: Cancel a specific client node session

Cancel the client node session with NODEP (session 3).

```
cancel session 3
```

Related commands

Table 1. Commands related to CANCEL SESSION

Command	Description
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
LOCK ADMIN	Prevents an administrator from accessing IBM Spectrum Protect.
LOCK NODE	Prevents a client from accessing the server.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.

CHECKIN LIBVOLUME (Check a storage volume into a library)

Use this command to add a sequential access storage volume or a cleaning tape to the server inventory for an automated library. The server cannot use a volume that physically resides in an automated library until that volume is checked in.

Important:

1. The CHECKIN LIBVOLUME command processing does not wait for a drive to become available, even if the drive is only in the IDLE state. If necessary, you can make a library drive available issuing the DISMOUNT VOLUME command to dismount the volume. After a library drive is available, reissue the CHECKIN LIBVOLUME command.
2. You do not define the drives, check in media, or label the volumes in an external library. The server provides an interface that external media management systems use to operate with the server.
3. When you check in WORM tapes other than 3592, you must use CHECKLABEL=YES or they are checked in as normal read/write tapes.

This command creates a background process that you can cancel with the CANCEL PROCESS command. To display information about background processes, use the QUERY PROCESS command.

For detailed and current drive and library support information, see the Supported Devices website for your operating system:

- **AIX** **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for SCSI libraries

```
>>-CHECKIn LIBVolume--library_name----->
. -SEARCH----No-.
>----+volume_name--+-----+----->
+-SEARCH----Yes--+-----+
|                '-| A |-|
'-SEARCH----Bulk--+-----+
|                '-| A |-|
. -OWNER-----"------
>--STATUS-----+PRIVate+-----+----->
+-SCRatch+    '-OWNER-----server_name-'
'-CLEaner-'

. -CHECKLabel----Yes----- . -SWAP----No-----
>----+-----+-----+----->
'-CHECKLabel-----+Yes-----+' '-SWAP-----+No-----+'
+-No-----+                '-Yes-'
'-Barcode-'

. -WAITTime----60----.
>--+-----+-----+-----+-----><
'-WAITTime----value-' '-CLEanings----number--'

A (SEARCH=Yes, SEARCH=Bulk)

|--+VOLRange--+-----volume_name1,volume_name2--+-----|
|                .-,-----|
|                V          |
'-VOLList--+-----volume_name+-----+'
'-FILE:--file_name-'
```

Syntax for 349X libraries

```
>>-CHECKIn LIBVolume--library_name----->
```

```

      .-SEARCH-----No-.
>-----+volume_name-----+-----+----->
      '-SEARCH-----Yes-----'
              '| A |-'

      .-OWNER-----"------
>--STATUS-----+PRIVATE-----+----->
      '-SCRATCH-' '-OWNER-----server_name-'

      .-CHECKLabel-----Yes-----
>-----+-----+-----+----->
      '-CHECKLabel-----+Yes+-' '-DEVType-----+3590+-'
              '-No--'              '-3592-'

      .-SWAP-----No----- .-WAITTime-----60-----
>-----+-----+-----+-----><
      '-SWAP-----+No--+-' '-WAITTime-----value-'
              '-Yes-'

```

A (SEARCH=Yes)

```

|---+VOLRange-----+volume_name1,volume_name2---+-----|
|           .-,----- .           |
|           V           |           |
'-VOLList-----+---+volume_name+---+-----'
              '-FILE:--file_name-'

```

Syntax for ACSLS libraries

```

>>-CHECKIn LIBVolume--library_name----->
      .-SEARCH-----No-.
>-----+volume_name-----+-----+----->
      '-SEARCH-----Yes-----'
              '| A |-'

      .-OWNER-----"------
>--STATUS-----+PRIVATE-----+----->
      '-SCRATCH-' '-OWNER-----server_name-'

      .-CHECKLabel-----Yes----- .-SWAP-----No-----
>-----+-----+-----+----->
      '-CHECKLabel-----+Yes+-' '-SWAP-----+No--+-'
              '-No--'              '-Yes-'

      .-WAITTime-----60-----
>-----+-----+-----+-----><
      '-WAITTime-----value-'

```

A (SEARCH=Yes)

```

|---+VOLRange-----+volume_name1,volume_name2---+-----|
|           .-,----- .           |
|           V           |           |
'-VOLList-----+---+volume_name+---+-----'
              '-FILE:--file_name-'

```

Parameters

library_name (Required)

Specifies the name of the library.

volume_name

Specifies the volume name of the storage volume that is being checked in. This parameter is required if SEARCH equals NO. Do not enter this parameter if the SEARCH parameter equals YES or BULK. If you are checking a volume into a SCSI library with multiple entry/exit ports, the volume in the lowest numbered slot is checked in.

STATUS (Required)

Specifies the volume status. Possible values are:

PRIVate

Specifies that the volume is a private volume that is mounted only when it is requested by name.

SCRatch

Specifies that the volume is a new scratch volume. This volume can be mounted to satisfy scratch mount requests during either data storage operations or export operations.

If a volume has an entry in volume history, you cannot check it in as a scratch volume.

CLEaner

Specifies that the volume is a cleaner cartridge and not a data cartridge. The CLEANINGS parameter is required for a cleaner cartridge and must be set to the number of cleaner uses.

CHECKLABEL=YES is not valid for checking in a cleaner cartridge. Use STATUS=CLEANER to check in a cleaner cartridge separately from a data cartridge.

OWNer

Specifies which library client owns a private volume in a library that is shared across a SAN. The volume for which you specify ownership must be a private volume. You cannot specify ownership for a scratch volume. Furthermore, you cannot specify an owner when you use SEARCH=YES or SEARCH=BULK.

When you issue the CHECKIN LIBVOLUME command, the server validates the owner. If you did not specify this parameter, then the server uses the default and delegates volume ownership to the owning library client, as recorded in the volume history file on the library manager. If the volume is not owned by any library client, then the server delegates ownership to the library manager.

SEARCH

Specifies whether the server searches the library to find volumes that were not checked in. This parameter is optional. The default is NO.

Possible values are:

No

Specifies that only the named volume is checked into the library.

For SCSI libraries: The server issues a request to have the volume inserted into a cartridge slot in the library or, if available, into an entry port. The cartridge slot or entry port is identified by its element address. **For 349X libraries:** The volume might already be in the library, or you can put it into the I/O station when prompted.

Yes

Specifies that the server searches the library for volumes to be checked in. You can use the VOLRANGE or VOLLIST parameter to limit the search. When you use this parameter, consider the following restrictions:

- If the library is shared between applications, the server might examine a volume that is required by another application. For 349X libraries, the server queries the library manager to determine all volumes that are assigned to the SCRATCH or PRIVATE category and to the INSERT category.
- For SCSI libraries, do not specify both SEARCH=YES and CHECKLABEL=NO in the same command.

Bulk

Specifies that the server searches the library's entry/exit ports for volumes that can be checked in automatically. This option applies to only SCSI libraries.

Important:

1. Do not specify both CHECKLABEL=NO and SEARCH=BULK.
2. You can use the VOLRANGE or VOLLIST parameter to limit the search.

VOLRange

Specifies a range of volume names that are separated by commas. You can use this parameter to limit the search for volumes to be checked in when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are within the specified range, the command completes without errors.

Specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
-----------	-------------

Parameter	Description
<code>volrange=bar110,bar130</code>	The 21 volumes are checked in: bar110, bar111, bar112,...bar129, bar130.
<code>volrange=bar11a,bar13a</code>	The 3 volumes are checked in: bar11a, bar12a, bar13a.
<code>volrange=123400,123410</code>	The 11 volumes are checked in: 123400, 123401, ...123409, 123410.

VOLLIST

Specifies a list of volumes. You can use this parameter to limit the search for volumes to be checked in when you specify `SEARCH=YES` (349X, ACSLS, and SCSI libraries) or `SEARCH=BULK` (SCSI libraries only). If there are no volumes in the library that are in the list, the command completes without errors.

Possible values are:

volume_name

Specifies one or more volumes names that are separated by commas and no intervening spaces. For example:
`VOLLIST=TAPE01,TAPE02.`

FILE: file_name

Specifies the name of a file that contains a list of volumes for the command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volumes TAPE01, TAPE02 and TAPE03, create a file, TAPEVOL, that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can specify the volumes for the command as follows: `VOLLIST=FILE:TAPEVOL.`

Attention: The file name is case-sensitive.

CHECKLabel

Specifies how or whether the server should read sequential media labels of volumes. This parameter is optional. The default is YES.

Possible values are:

Yes

Specifies that an attempt is made to read the media label during check-in.

Attention:

1. For SCSI libraries, do not specify both `SEARCH=YES` and `CHECKLABEL=NO` in the same command.
2. For WORM media other than 3592, you must specify YES.

No

Specifies that the media label is not read during check-in. However, suppressing label checking can result in future errors (for example, either a wrong label or an improperly labeled volume can cause an error). For 349X and ACSLS libraries, specify NO to avoid loading cartridges into a drive to read the media label. These libraries always return the external label information about cartridges, and IBM Spectrum Protect™ uses that information.

Barcode

Specifies that the server reads the bar code label if the library has a bar code reader and the volumes have external bar code labels. You can decrease the check-in time by using the bar code. This parameter applies only to SCSI libraries.

If the bar code reader cannot read the bar code label, or if the tape does not have a bar code label, the server mounts the tape and reads the internal label.

DEVType

Specifies the device type for the volume that is being checked in. This parameter is required if none of the drives in this library have defined paths.

3590

Specifies that the device type for the volume that is being checked in is 3590.

3592

Specifies that the device type for the volume that is being checked in is 3592.

SWAP

Specifies whether the server swaps volumes if an empty library slot is not available. The volume that is selected for the swap operation (target swap volume) is ejected from the library and replaced with the volume that is being checked in. The server identifies a target swap volume by checking for an available scratch volume. If none exists, the server identifies the least frequently mounted volume.

This parameter is optional. The default is NO. This parameter applies only if there is a volume name that is specified in the command. Possible values are:

No

Specifies that the server checks in the volume only if an empty slot is available.

Yes

Specifies that if an empty slot is not available, the server swaps cartridges to check in the volume.

WAITTime

Specifies the number of minutes that the server waits for you to reply or respond to a request. Specify a value in the range 0-9999. If you want to be prompted by the server, specify a wait time greater than zero. The default value is 60 minutes. For example, suppose the server prompts you to insert a tape into the entry/exit port of a library. If you specified a wait time of 60 minutes, the server issues a request and waits 60 minutes for you to reply. Suppose, on the other hand, you specify a wait time of 0. If you already inserted a tape, a wait time of zero causes the operation to continue without prompting. If you have *not* inserted a tape, a wait time of zero will cause the operation to fail.

CLEanings

Enter the recommended value for the individual cleaner cartridge (usually indicated on the cartridge). Cleanings apply only to SCSI libraries. This parameter is required if STATUS=CLEANER.

If more than one cleaner is checked into the library, only one is used until its CLEANINGS value decreases to zero. Another cleaner is then selected, and the first cleaner can be checked out and discarded.

Example: Check a volume into a SCSI library

Check in a volume named `WPDV00` into the SCSI library named `AUTO`.

```
checkin libvolume auto wpdv00 status=scratch
```

Example: Use a bar code reader to scan a library for a cleaner cartridge

Scan a SCSI library named `AUTOLIB1` and, using the bar code reader, look for cleaner cartridge `CLNV`. Use `SEARCH=YES`, but limit the search by using the `VOLLIST` parameter.

```
checkin libvolume autolib1 search=yes vollist=cleanv status=cleaner  
cleanings=10 checklabel=barcode
```

Example: Scan a library to put unused volumes in a specific range in scratch status

Scan a 349X library named `ABC`, and limit the search to a range of unused volumes `BAR110` to `BAR130` and put them in scratch status.

```
checkin libvolume abc search=yes volrange=bar110,bar130  
status=scratch
```

Example: Scan a library to put a specific volume in scratch status

Use the barcode reader to scan a SCSI library named `MYLIB` for `VOL1`, and put it in scratch status.

```
checkin libvolume mylib search=yes vollist=voll status=scratch  
checklabel=barcode
```

Related commands

Table 1. Commands related to CHECKIN LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.

Command	Description
CANCEL PROCESS	Cancels a background server process.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
UPDATE LIBVOLUME	Changes the status of a storage volume.

CHECKOUT LIBVOLUME (Check a storage volume out of a library)

Use this command to remove a sequential access storage volume from the server inventory for an automated library. This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Restrictions:

1. Check out processing does not wait for a drive to become available, even if the drive is in the IDLE state. If necessary, you can make a library drive available by dismounting the volume with the DISMOUNT VOLUME command. After a drive is available, the CHECKOUT LIBVOLUME command can be reissued.
2. Before checking out volumes from a 349X library, ensure that the 349x Cartridge Input and Output facility has enough empty slots for the volumes to be checked out. The 3494 Library Manager does not inform an application that the Cartridge Input and Output facility is full. It accepts requests to eject a cartridge and waits until the Cartridge Input and Output facility is emptied before returning to the server. IBM Spectrum Protect™ might appear to be hung when it is not. Check the library and clear any intervention requests.
3. Before checking volumes out of an ACSLS library, ensure that the CAP priority in ACSLS is greater than zero. If the CAP priority is zero, then you must specify a value for the CAP parameter on the CHECKOUT LIBVOLUME command.

For detailed and current drive and library support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for SCSI library

```

>>-CHECKOut LIBVolume--library_name----+volume_name+----->
                                     '-| A |-----'

.-REMove----Bulk-----.-CHECKLabel----Yes-----
>+-----+-----+-----+-----+-----+----->
  '-REMove----+Yes---+'  '-CHECKLabel----+Yes---+'
      +-No---+                '-No--'
      '-Bulk-'

.-FORCE----No-----
>+-----+-----+-----+-----+-----+----->>
  '-FORCE----+No---+'
      '-Yes-'

A

|--+VOLRange------volume_name1,volume_name2--+-----|
|          .-,-----|
|          V          |
|'-VOLList--++---volume_name+---+-----|
|          '-FILE:--file_name-'

```

Syntax for 349X library

```

>>-CHECKOut LIBVolume--library_name----+volume_name+----->
                                     '-| A |-----'

.-REMove----Bulk-----
>+-----+-----+-----+-----+-----+-----><
  '-REMove----+Yes---+'
      +-No---+
      '-Bulk-'

A

|--+VOLRange------volume_name1,volume_name2--+-----|
|          .-,-----|
|          V          |
|'-VOLList--++---volume_name+---+-----|
|          '-FILE:--file_name-'

```

Syntax for ACSLS library

```

>>-CHECKOut LIBVolume--library_name----+volume_name+----->
                                     '-| A |-----'

.-REMove----Yes-----
>+-----+-----+-----+-----+-----+-----><
  '-REMove----+Yes---+'  '-CAP-----x,y,z---+'
      +-No---+
      '-Bulk-'

A

|--+VOLRange------volume_name1,volume_name2--+-----|
|          .-,-----|
|          V          |
|'-VOLList--++---volume_name+---+-----|
|          '-FILE:--file_name-'

```

Parameters

library_name (Required)
 Specifies the name of the library.

volume_name

Specifies the volume name.

VOLRange

Specifies two volume names separated by a comma. This parameter is a range of volumes to be checked out. If there are no volumes in the library that are within the specified range, the command completes without errors.

Specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
<code>volrange=bar110,bar130</code>	The 21 volumes are checked out: bar110, bar111, bar112,...bar129, bar130.
<code>volrange=bar11a,bar13a</code>	The 3 volumes are checked out: bar11a, bar12a, bar13a.
<code>volrange=123400,123410</code>	The 11 volumes are checked out: 123400, 123401, ...123409, 123410.

VOLList

Specifies a list of volumes to check out. If there are no volumes in the library that are in the list, the command completes without errors.

Possible values are:

volume_name

Specifies the names of one or more values that are used for the command. Example: `VOLLIST=TAPE01,TAPE02`.

FILE:file_name

Specifies the name of a file that contains a list of volumes for the command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volumes TAPE01, TAPE02 and TAPE03, create a file, TAPEVOL, that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can specify the volumes for the command as follows: `VOLLIST=FILE:TAPEVOL`.

Attention: The file name is case-sensitive.

REMOve

Specifies that the server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values, depending on the type of library, are YES, BULK, and NO. The response of the server to each of those options and the default values are described in the following sections.

349X libraries: The default is BULK. The following table shows how the server responds for 349X libraries.

Table 1. How the server responds for 349X libraries

REMOVE=YES	REMOVE=BULK	REMOVE=NO
The 3494 Library Manager ejects the cartridge to the convenience I/O station.	The 3494 Library Manager ejects the cartridge to the high-capacity output facility.	The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the INSERT category for use by other applications.

SCSI libraries: The default is BULK. The following table shows how the server responds for a SCSI libraries.

Table 2. How the server responds for SCSI libraries

If a library . . .	And REMOVE=YES, then...	And REMOVE=BULK, then...	And REMOVE=NO, then...

If a library . . .	And REMOVE=YES, then...	And REMOVE=BULK, then...	And REMOVE=NO, then...
<i>Does not have entry/exit ports</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has entry/exit ports and an entry/exit port is available</i>	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has entry/exit ports, but no ports are available</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server waits for an entry/exit port to be made available.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.

ACSLs libraries: The default is YES. If the parameter is set to YES, and the cartridge access port (CAP) has an automatic selection priority value of 0, you must specify a CAP ID. The following table shows how the server responds for ACSLS libraries.

Table 3. How the server responds for ACSLS libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station, and deletes the volume entry from the server library inventory.	The server does not eject the cartridge. The server deletes the volume entry from the server library inventory and leaves the volume in the library.

CHECKLabel

Specifies how or whether the server reads sequential media labels of volumes.

Attention: This parameter does not apply to IBM® 349X or ACSLS libraries.

This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the server attempts to read the media label to verify that the correct volume is being checked out.

No

Specifies that during checkout the media label is not read. This improves performance because the read process does not occur.

FORCE

Specifies whether the server checks out a volume if an input/output (I/O) error occurs when reading the label.

Attention: This parameter does not apply to IBM 349X or ACSLS libraries.

This parameter is optional. The default is NO. Possible values are:

- No
The server does not check out a storage volume if an I/O error occurs when reading the label.
- Yes
The server checks out the storage volume even if an I/O error occurs.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (x,y,z), issue the QUERY CAP command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

- x
The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.
- y
The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.
- z
The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Example: Check out a volume and check the label

Check out the volume that is named EXB004 from the library named FOREST. Read the label to verify the volume name, but do not move the volume out of the library.

```
checkout libvolume forest exb004 checklabel=yes remove=no
```

Related commands

Table 4. Commands related to CHECKOUT LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CANCEL PROCESS	Cancels a background server process.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
UPDATE LIBVOLUME	Changes the status of a storage volume.

CLEAN DRIVE (Clean a drive)

Use this command when you want IBM Spectrum Protect™ to immediately load a cleaner cartridge into a drive regardless of the cleaning frequency.

There are special considerations if you plan to use this command with a SCSI library that provides automatic drive cleaning through its device hardware.

Restriction: You cannot run the CLEAN DRIVE command for a drive whose only path source is a NAS file server.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-CLEAN DRIVE--library_name--drive_name-----<<
```

Parameters

library_name (Required)
Specifies the name of the library to which the drive is assigned.

drive_name (Required)
Specifies the name of the drive.

Example: Clean a specific tape drive

You have already defined a library named AUTOLIB by using the DEFINE LIBRARY command, and you have already checked a cleaner cartridge into the library using the CHECKIN LIBVOL command. Inform the server that TAPEDRIVE3 in this library requires cleaning.

```
clean drive autolib tapedrive3
```

Related commands

Table 1. Commands related to CLEAN DRIVE

Command	Description
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DRIVE	Deletes a drive from a library.
QUERY DRIVE	Displays information about drives.
UPDATE DRIVE	Changes the attributes of a drive.

COMMIT (Control committing of commands in a macro)

Use this command to control when a command is committed in a macro and to update the database when commands complete processing. When issued from the console mode of the administrative client, this command does not generate a message.

If an error occurs while processing the commands in a macro, the server stops processing the macro and rolls back any changes (since the last COMMIT). After a command is committed, it cannot be rolled back.

Ensure that your administrative client session is not running with the ITEMCOMMIT option if you want to control command processing. The ITEMCOMMIT option commits commands inside a script or a macro as *each* command is processed.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-COMMIT-----<<
```

Parameters

None.

Example: Control committing of commands in a macro

From the interactive mode of the administrative client, register and grant authority to new administrators using a macro named REG.ADM. Changes are committed after each administrator is registered and is granted authority.

Macro Contents:

```
/* REG.ADM-register policy admin & grant authority*/
REGister Admin sara hobby
GRant AUTHority sara CLasses=Policy
COMMIT /* Commits changes */
REGister Admin ken plane
GRant AUTHority ken CLasses=Policy
COMMIT /* Commits changes */
```

Command

```
macro reg.adm
```

Related commands

Table 1. Commands related to COMMIT

Command	Description
MACRO	Runs a specified macro file.
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.

Related concepts:

Administrative client macros



CONVERT STGPOOL (Convert a storage pool to a container storage pool)

Use this command to convert a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL) to a directory-container or a cloud-container storage pool. You can use container storage pools for both inline and client-side data deduplication.

Restrictions: The following restrictions apply to storage pool conversion:

- You can convert a storage pool only once.
- You cannot update the storage pool during conversion processing. Migration and data movement processes are unavailable.
- You must update all policies to ensure that the destination specifies a storage pool that is not converted or undergoing conversion.

During conversion processing, all data from the source storage pool is moved to the target storage pool. When the process is completed, the source storage pool becomes unavailable. When a storage pool is unavailable, you are unable to write any data to it. The source storage pool is eligible for deletion but is not automatically deleted. You can restore data from the source storage pool if necessary.

Attention: During storage pool conversion, data is deleted from copy storage pools and active-data storage pools. This action occurs even if you specified the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool.

Privilege class

To issue this command, you must have restricted storage privilege.

Syntax

```
>>-CONvert STGpool--source_stgpool--target_stgpool----->
```

```

.-MAXPRocess-----8-----
>-----+-----+-----+-----+-----<
'-MAXPRocess-----number---' '-DURation-----minutes-'

```

Parameters

source_stgpool (Required)

Specify a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL) for backup and archive processing. This parameter is required.

target_stgpool (Required)

Specify the name of an existing directory-container or cloud-container storage pool that the storage pool is converted to. This parameter is required the first time that you issue this command.

Tip: If you restart storage pool conversion and the target storage pool is different than the value that is specified the first time that you issued the CONVERT STGPOOL command, the command fails.

MAXPRocess

Specifies the maximum number of parallel processes that can be used to convert data in the storage pool. This parameter is optional. You can specify a number in the range 1 - 99. The default value is 8.

Tip: Changes to the default value are automatically saved. If you restart storage pool conversion and the parameter value is different than the value that is specified the first time that you issued the CONVERT STGPOOL command, the most recently specified value is used.

DURation

Specifies the maximum number of minutes that a conversion should take before it is canceled. When the specified number of minutes elapses, the server cancels all conversion processes for the storage pool. You can specify a number in the range 1 - 9999. This parameter is optional. If you do not specify this parameter, the conversion runs until it is completed.

Tip: Storage pool conversion for large storage pools can take days to complete. Use this parameter to limit the amount of time for storage pool conversion daily. As a best practice, schedule conversion for at least 2 hours for a storage pool that uses a FILE type device class and at least 4 hours for VTL.

Example: Convert a storage pool and specify a maximum number of processes

Convert a storage pool that is named DEDUPPOOL1, move the data to a container storage pool that is named DIRPOOL1, and specify 25 maximum processes.

```
convert stgpool deduppool1 dirpool1 maxprocess=25
```

Table 1. Commands related to CONVERT STGPOOL

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY CLEANUP	Query the cleanup status of a source storage pool.
QUERY CONVERSION	Query conversion status of a storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
REMOVE DAMAGED	Removes damaged data from a source storage pool.

COPY commands

Use the COPY commands to create a copy of IBM Spectrum Protect™ objects or data.

- COPY ACTIVEDATA (Copy active backup data from a primary storage pool to an active-data pool)
- COPY CLOPTSET (Copy a client option set)
- COPY DOMAIN (Copy a policy domain)
- COPY MGMTCLASS (Copy a management class)
- COPY POLICYSET (Copy a policy set)
- COPY PROFILE (Copy a profile)
- COPY SCHEDULE (Copy a client or an administrative command schedule)
- COPY SCRIPT (Copy an IBM Spectrum Protect script)
- COPY SERVERGROUP (Copy a server group)

COPY ACTIVEdata (Copy active backup data from a primary storage pool to an active-data pool)

Use this command to copy active versions of backup data from a primary storage pool to an active-data pool. The primary benefit of active-data pools is fast client restores. Copy your active data regularly to ensure that the data is protected in case of a disaster.

If a file already exists in the active-data pool, the file is not copied unless the copy of the file in the active-data pool is marked damaged. However, a new copy is not created if the file in the primary storage pool is also marked damaged. In a random-access storage pool, neither cached copies of migrated files nor damaged primary files are copied.

If migration for a storage pool starts while active data is being copied, some files might be migrated before they are copied. For this reason, you should copy active data from storage pools that are higher in the migration hierarchy before copying active data from storage pools that are lower. Be sure a copy process is complete before beginning another.

Remember:

- You can only copy active data from storage pools that have a data format of NATIVE or NONBLOCK.
- Issuing this command for a primary storage pool that is set up for data deduplication removes duplicate data, if the active-data pool is also set up for data deduplication.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the active-data pool from which active versions of backup data are being copied.

Syntax

```
>>-COPY ACTIVEdata--primary_pool_name--active-data_pool_name---->
. -MAXProcess-----1-----
>--+-----+----->
' -MAXProcess-----number--- '

. -Preview-----No----- . -Wait-----No-----
>--+-----+-----+----->
' -Preview-----+No-----+ ' ' -Wait-----+No-----+ '
      +-Yes-----+          ' -Yes- '
      |               (1) |
      '-VOLUMESONLY-----'

. -SHREDTONOshred-----No-----
>--+-----+----->>
' -SHREDTONOshred-----+No-----+ '
      ' -Yes- '

```

Notes:

1. The VOLUMESONLY parameter applies to sequential-access storage pools only.

Parameters

primary_pool_name (Required)

Specifies the primary storage pool.

active_data_pool_name (Required)

Specifies the active-data pool.

MAXProcess

Specifies the maximum number of parallel processes to use for copying files. This parameter is optional. Enter a value from 1 to 999. The default is 1.

Using multiple, parallel processes may improve throughput for the COPY ACTIVEdata command. The expectation is that the time needed to copy active data will be decreased by using multiple processes. However, when multiple processes are running, in some cases one or more of the processes might need to wait to use a volume that is already in use by a different COPY ACTIVEdata process.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential-access volume, the server uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other server and system activity, and also on the mount limits of the device classes for the sequential-access storage pools that are involved when copying active data.

Each process needs a mount point for active-data pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are copying active data from a sequential-access storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device type is not FILE, an additional drive. For example, suppose you specify a maximum of 3 processes to copy a primary sequential storage pool to an active-data pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least six, and at least six mount points and six drives must be available.

To use the PREVIEW parameter, only one process is used, and no mount points or drives are needed.

Preview

Specifies whether you want to preview but not actually copy any active data. The preview displays the number of files and bytes to be copied and a list of the primary storage pool volumes that you must mount. This parameter is optional. The default is NO. Possible values are:

No

Specifies that active data will be copied.

Yes

Specifies that you want to preview the process but not copy any data.

VOLumesonly

Specifies that you want to preview the process only as a list of the volumes that must be mounted. This choice requires the least processing time.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files may have already been copied prior to the cancellation.

Yes

Specifies that the server performs this operation in the foreground. You must wait for the operation to complete before continuing with other tasks. The server displays the output messages to the administrative client when the operation completes.

You cannot specify WAIT=YES from the server console.

SHREDTONOshred

Specifies whether data should be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the server does not allow data to be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. If the primary storage pool enforces shredding and the active-data pool does not, the operation will fail.

Yes

Specifies that the server does allow data to be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. The data in the active-data pool will not be shredded when it is deleted.

Example: Copy primary storage pool data to active-data pool

Copy the active data from a primary storage pool named PRIMARY_POOL to the active-data pool named ACTIVEPOOL. Issue the command:

Related commands

Table 1. Commands related to COPY ACTIVEDATA

Command	Description
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY DOMAIN	Displays information about policy domains.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE DOMAIN	Changes the attributes of a policy domain.
UPDATE STGPOOL	Changes the attributes of a storage pool.

COPY CLOPTSET (Copy a client option set)

Use this command to copy a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-COPY CLOptset--current_option_set_name--new_option_set_name-><
```

Parameters

current_option_set_name (Required)

Specifies the name of the client option set to be copied.

new_option_set_name (Required)

Specifies the name of the new client option set. The maximum length of the name is 64 characters.

Example: Copy a client option set

Copy a client option set named ENG to a new client option set named ENG2.

```
copy cloptset eng eng2
```

Related commands

Table 1. Commands related to COPY CLOPTSET

Command	Description
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

COPY DOMAIN (Copy a policy domain)

Use this command to create a copy of a policy domain.

The server copies the following information to the new domain:

- Policy domain description
- Policy sets in the policy domain (including the ACTIVE policy set, if a policy set is activated)
- Management classes in each policy set (including the default management class, if assigned)
- Copy groups in each management class

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-COPY Domain--current_domain_name--new_domain_name-----<<
```

Parameters

current_domain_name (Required)

Specifies the policy domain to copy.

new_domain_name (Required)

Specifies the name of the new policy domain. The maximum length of this name is 30 characters.

Example: Copy a policy domain to a new policy domain

Copy the STANDARD policy domain to a new policy domain, ENGPOLDOM, by entering the following command:

```
copy domain standard engpoldom
```

ENGPOLDOM now contains the standard policy set, management class, backup copy group, and archive copy group.

Related commands

Table 1. Commands related to COPY DOMAIN

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE MGMTCLASS	Defines a management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY DOMAIN	Displays information about policy domains.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE DOMAIN	Changes the attributes of a policy domain.
UPDATE MGMTCLASS	Changes the attributes of a management class.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

COPY MGMTCLASS (Copy a management class)

Use this command to create a copy of a management class within the same policy set.

The server copies the following information to the new management class:

- Management class description
- Copy groups defined to the management class
- Any attributes for managing files for IBM Spectrum Protect™ for Space Management clients

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new management class belongs.

Syntax

```
>>-COpy MGmtclass--domain_name--policy_set_name----->
>>--current_class_name--new_class_name-----<
```

Parameters

domain_name (Required)

Specifies the policy domain to which the management class belongs.

policy_set_name (Required)

Specifies the policy set to which the management class belongs.

current_class_name (Required)

Specifies the management class to copy.

new_class_name (Required)

Specifies the name of the new management class. The maximum length of this name is 30 characters.

Example: Copy a management class to a new management class

Copy the management class ACTIVEFILES to a new management class, FILEHISTORY. The management class is in policy set VACATION in the EMPLOYEE_RECORDS policy domain.

```
copy mgmtclass employee_records vacation
activefiles filehistory
```

Related commands

Table 1. Commands related to COPY MGMTCLASS

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.

COPY POLICYSET (Copy a policy set)

Use this command to copy a policy set (including the ACTIVE policy set) within the same policy domain.

The server copies the following information to the new policy set:

- Policy set description
- Management classes in the policy set (including the default management class, if assigned)
- Copy groups in each management class

The policies in the new policy set do not take effect unless you make the new set the ACTIVE policy set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new policy set belongs.

Syntax

```
>>-COPY Policyset--domain_name--current_set_name--new_set_name-><
```

Parameters

domain_name (Required)

Specifies the policy domain to which the policy set belongs.

current_set_name (Required)

Specifies the policy set to copy.

new_set_name (Required)

Specifies the name of the new policy set. The maximum length of this name is 30 characters.

Example: Copy a policy set to a new policy set

Copy the policy set `VACATION` to the new policy set `HOLIDAY` in the `EMPLOYEE_RECORDS` policy domain.

```
copy policyset employee_records vacation holiday
```

Related commands

Table 1. Commands related to COPY POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

COPY PROFILE (Copy a profile)

Use this command on a configuration manager to copy a profile and all its associated object names to a new profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-COpy PROFIle--current_profile_name--new_profile_name-----><
```

Parameters

current_profile_name (Required)

Specifies the profile to copy.

new_profile_name (Required)

Specifies the name of the new profile. The maximum length of the profile name is 30 characters.

Example: Make a copy of a profile

Copy a profile named `VAL` to a new profile named `VAL2`.

```
copy profile val val2
```

Related commands

Table 1. Commands related to COPY PROFILE

Command	Description
---------	-------------

Command	Description
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

COPY SCHEDULE (Copy a client or an administrative command schedule)

Use this command to create a copy of a schedule.

The COPY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Table 1. Commands related to COPY SCHEDULE

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY SCHEDULE	Displays information about schedules.
UPDATE SCHEDULE	Changes the attributes of a schedule.

- **COPY SCHEDULE (Create a copy of a schedule for client operations)**
Use the COPY SCHEDULE command to create a copy of a schedule for client operations. You can copy a schedule within a policy domain or from one policy domain to another policy domain. Use the DEFINE ASSOCIATION command to associate the new schedule with the client nodes.
- **COPY SCHEDULE (Create a copy of a schedule for administrative operations)**
Use the COPY SCHEDULE command to create a copy of an administrative command schedule.

COPY SCHEDULE (Create a copy of a schedule for client operations)

Use the COPY SCHEDULE command to create a copy of a schedule for client operations. You can copy a schedule within a policy domain or from one policy domain to another policy domain. Use the DEFINE ASSOCIATION command to associate the new schedule with the client nodes.

Privilege class

To copy a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which you are copying the schedule.

Syntax

```
>>-COpy SChedule--current_domain_name--current_sched_name----->
                                     .-current_sched_name-.
>>-new_domain_name-+-----+----->
                                     '-new_sched_name-----'

                                     .-REPlace----No----- .
>--+-----+----->>
                                     '-REPlace----+No--+-'
                                     '-Yes-'
```

Parameters

current_domain_name (Required)

Specifies the name of the policy domain that contains the schedule you want to copy.

current_sched_name (Required)

Specifies the name of the schedule you want to copy.

new_domain_name (Required)

Specifies the name of a policy domain to which you want to copy the new schedule.

new_sched_name

Specifies the name of the new schedule. You can specify up to 30 characters for the name.

If you do not specify this name, the name of the original schedule is used.

If the schedule name is already defined in the policy domain, you must specify REPLACE=YES, or the command fails.

REPlace

Specifies whether to replace a client schedule. The default is NO. The values are:

No

Specifies that a client schedule is not replaced.

Yes

Specifies that a client schedule is replaced.

Example: Copy a schedule from one policy domain to another

Copy the WEEKLY_BACKUP schedule that belongs to policy domain EMPLOYEE_RECORDS to the PROG1 policy domain and name the new schedule WEEKLY_BACK2. If there is already a schedule with this name defined in the PROG1 policy domain, do not replace it.

```
copy schedule employee_records weekly_backup
prog1 weekly_back2
```

COPY SCHEDULE (Create a copy of a schedule for administrative operations)

Use the COPY SCHEDULE command to create a copy of an administrative command schedule.

Privilege class

To copy an administrative command schedule, you must have system privilege.

Syntax

```
>>-COpy SChedule--current_sched_name--new_sched_name----->
```

```

>--Type---Administrative--+-REPlace---No-----><
                              '-REPlace---No---+'
                              '-Yes-'

```

Parameters

current_schedule_name (Required)

Specifies the name of the schedule you want to copy.

new_schedule_name (Required)

Specifies the name of the new schedule. You can specify up to 30 characters for the name.

If the schedule name is already defined, you must specify REPLACE=YES, or the command fails.

Type=Administrative

Specifies that an administrative command schedule is to be copied.

REPlace

Specifies whether to replace an administrative command schedule. The default is NO. The values are:

No

Specifies that an administrative command schedule is not replaced.

Yes

Specifies that an administrative command schedule is replaced.

Example: Copy an administrative command schedule to another schedule

Copy the administrative command schedule, DATA_BACKUP and name the schedule DATA_ENG. If there is already a schedule with this name, replace it.

```
copy schedule data_backup data_eng
type=administrative replace=yes
```

COPY SCRIPT (Copy an IBM Spectrum Protect script)

Use this command to copy an existing IBM Spectrum Protect™ script to a new script with a different name.

Privilege class

To issue this command, you must have operator, policy, storage, or system privilege.

Syntax

```
>>-COpy SCRipt--current_script_name--new_script_name -----><
```

Parameters

current_script_name (Required)

Specifies the name of the script you want to copy.

new_script_name (Required)

Specifies the name of the new script. You can specify up to 30 characters for the name.

Example: Make a copy of a script

Copy script TESTDEV to a new script and name it ENGDEV.

```
copy script testdev engdev
```

Related commands

Table 1. Commands related to COPY SCRIPT

Command	Description
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

COPY SERVERGROUP (Copy a server group)

Use this command to create a copy of a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-COPY SERVERGroup--current_group_name--new_group_name-----><
```

Parameters

current_group_name (Required)

Specifies the server group to copy.

new_group_name (Required)

Specifies the name of the new server group. The maximum length of this name is 64 characters.

Example: Make a copy of a server group

Copy the server group GRP_PAYROLL to the new group HQ_PAYROLL.

```
copy servergroup grp_payroll hq_payroll
```

Related commands

Table 1. Commands related to COPY SERVERGROUP

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVER	Defines a server for server-to-server communications.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVER	Deletes the definition of a server.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVER	Updates information about a server.
UPDATE SERVERGROUP	Updates a server group.

DEACTIVATE DATA (Deactivate data for a client node)

Use this command to specify that active data that was backed up for an application client node before a specified date is no longer needed. The command marks the data as inactive so it can be deleted according to your data retention policies.

Restriction: The DEACTIVATE DATA command applies only to application clients that protect Oracle databases.

When you issue the DEACTIVATE DATA command, all active backup data that was stored before the specified date becomes inactive. The data can no longer be retrieved, and is deleted when it expires.

The DEACTIVATE DATA command affects only the files that were copied to the server before the specified date and time. Files that were copied after the specified date are still accessible, and the client can still access the server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEACTivate DAta--node_name--TODate---date----->
      .-Totime---23:59:59-.  .-Wait---No-----
>--+-----+-----+-----+----->>
      '-Totime---time-----'  '-Wait---+No---+'
                                   '-Yes-'
```

Parameters

node_name (Required)

Specifies the name of an application client node whose data is to be deactivated.

TODate (Required)

Specifies the date to use to select the backup files to deactivate. IBM Spectrum Protect™ deactivates only those files with a date on or before the date you specify. You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	01/23/2014
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-30 or -30. To deactivate files that are 30 or more days old, you can specify TODAY-30 or -30.
EOLM	End of last month. The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To deactivate files that were active a day before the last day of the previous month.
BOTM	Beginning of this month. The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To deactivate files that were active on the 10th day of the current month.

TOTime

Specifies that you want to deactivate files that were created on the server before this time on the specified date. This parameter is optional. The default is the end of the day (23:59:59). Specify the time by using one of the following values:

Value	Description	Example
-------	-------------	---------

Value	Description	Example
HH:MM:SS	A specific time on the specified date	12:30:22
NOW	The current time on the specified date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified date	NOW+03:00 or +03:00. If you issue the DEACTIVATE DATA command at 9:00 with TOTIME=NOW+03:00 or TOTIME=+03:00, IBM Spectrum Protect deactivates files that were put on the server at 12:00 or earlier on the specified date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified date	NOW-03:30 or -03:30. If you issue the DEACTIVATE DATA command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, IBM Spectrum Protect deactivates files that were put on the server at 5:30 or earlier on the specified date.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. Specify the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages that are related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must complete before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Deactivate data for a data protection client node

The client node BANDIT is an IBM Spectrum Protect for Databases: Data Protection for Oracle application client. All of the backup data is active, and so all of the backup data is retained. The following command deactivates data that was backed up before January 3, 2014, so it can be deleted when it expires.

```
deactivate data bandit todate=01/23/2014
```

To periodically deactivate data so it can be deleted when it expires, you might run the following command from within a client schedule.

```
deactivate data bandit todate=today
```

Related commands

Table 1. Commands related to DEACTIVATE DATA

Command	Description
DECOMMISSION NODE	Decommissions an application or system.
DECOMMISSION VM	Decommissions a virtual machine.

DECOMMISSION commands

Use the DECOMMISSION commands to remove client nodes from the production environment. Client nodes include applications, systems, and virtual machines.

- DECOMMISSION NODE (Decommission an application or system)
- DECOMMISSION VM (Decommission a virtual machine)

DECOMMISSION NODE (Decommission an application or system)

Use this command to remove an application or system client node from the production environment. Any backup data that is stored for the client node expires according to policy settings unless you explicitly delete the data.

Attention: This action cannot be reversed and causes deletion of data. Although this command does not delete the client node definition until after its data expires, you cannot recommission the client node. After you issue this command, the client node cannot access the server and its data is not backed up. The client node is locked, and can be unlocked only to restore files. File spaces that belong to the client node, and the client node itself, are eventually removed.

By using this command, you can decommission the following types of client nodes:

Application client nodes

Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:

- IBM Spectrum Protect™ Snapshot
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Virtual Environments

System client nodes

System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

When a client node is no longer needed in the production environment, you can issue this command to initiate a gradual, controlled decommission operation. The command completes the following actions:

- Deletes all schedule associations for the client node. Schedules are no longer run on the client node. This action is equivalent to issuing the DELETE ASSOCIATION command for every schedule with which the client node is associated.
- Prevents the client from accessing the server. This action is equivalent to issuing the LOCK NODE command.

After the command finishes, client node data is no longer backed up to the server. Data that was backed up before the client node was decommissioned is not immediately deleted from the server. However, all backup file versions, including the most recent backup, are now inactive copies. The client files are retained on the server according to your storage management policies.

After all data retention periods expire, and all client backup and archive file copies are removed from server storage, IBM Spectrum Protect deletes the file spaces that belong to the decommissioned node. This action is equivalent to issuing the DELETE FILESPACE command.

After the file spaces for the decommissioned node are deleted, the node definition is deleted from the server. This action is equivalent to issuing the REMOVE NODE command.

After you decommission a client node, but before it is removed from the server, you can use the QUERY NODE command to verify that the client node is decommissioned.

Restriction: You cannot decommission a client node that is configured for replication. You can determine a client node's replication state by using the QUERY NODE command. If a client node is configured for replication, you can remove the client node from replication by using the REMOVE REPLNODE command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DECommission Node--node_name--+-----+-----><
                               .-Wait-----No-----
                               '-Wait-----+No--+-'
                               '-Yes-'
```

Parameters

node_name (Required)

Specifies the name of the client node to be decommissioned.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. You can specify the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages that are related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must complete before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Decommission a client node

Decommission the client node CODY.

```
decommission node cody
```

Related commands

Table 1. Commands related to DECOMMISSION NODE

Command	Description
DECOMMISSION VM	Decommissions a virtual machine.
DEACTIVATE DATA	Deactivates data for a client node.

DECOMMISSION VM (Decommission a virtual machine)

Use this command to remove an individual virtual machine within a data center node. The file space that represents the virtual machine is deleted from the server only after its backup data expires.

Attention: This command cannot be reversed and causes deletion of data. Although this command does not delete the virtual machine file space until after its data expires, you cannot recommission the virtual machine.

When a virtual machine is no longer needed in your production environment, you can issue this command to initiate a staged removal of the virtual machine file space from the server. The DECOMMISSION VM command marks all data that was backed up for the virtual machine as inactive, so it can be deleted according to your data retention policies. After all data that was backed up for the virtual machine expires, the file space that represents the virtual machine is deleted. The DECOMMISSION VM command affects only the virtual machine that you identify. The data center node, and the other virtual machines that are hosted by the data center node are not affected.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEComMission VM--node_name--vm_name----->
                                     .-Wait----No-----
>--+-----+-----+-----+-----><
'-NAMEType--FSID--' '-Wait----+Yes--+'
                                     '-No--'
```

Parameters

node_name (Required)

Specifies the name of the data center node that hosts the virtual machine to be decommissioned.

vm_name (Required)

Identifies the file space that represents the virtual machine to be decommissioned. Each virtual machine that is hosted by a data center node is represented as a file space.

If the name includes one or more spaces, you must enclose the name in double quotation marks when you issue the command.

By default, the server interprets the file space name that you enter by using the server code page and also attempts to convert the file space name from the server code page to the UTF-8 code page. Conversion might fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

If the name of the virtual machine is a non-English-language name, this parameter must specify the file space ID (FSID). By specifying the NAMETYPE parameter, you can instruct the server to interpret the file space name by its file space ID (FSID) instead.

NAMETYPE

Specify how you want the server to interpret the file space name that you enter to identify the virtual machine. This parameter is useful when the server has clients with Unicode support. You can specify the following value:

FSID

The server interprets the file space name by its file space ID (FSID).

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. You can specify the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages that are related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must complete before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Examples: Decommission a virtual machine

Decommission the virtual machine CODY.

```
decommission vm dept06node cody
```

Decommission the virtual machine CODY 2.

```
decommission vm dept06node "cody 2"
```

Decommission a virtual machine by specifying its file space ID.

```
decommission vm dept06node 7 nametype=fsid
```

Related commands

Table 1. Commands related to DECOMMISSION VM

Command	Description
DECOMMISSION NODE	Decommissions an application or system.
DEACTIVATE DATA	Deactivates data for a client node.

DEFINE commands

Use the DEFINE commands to create IBM Spectrum Protect™ objects.

- DEFINE ALERTTRIGGER (Define an alert trigger)
- DEFINE ASSOCIATION (Associate client nodes with a schedule)
- DEFINE BACKUPSET (Define a backup set)
- DEFINE CLIENTACTION (Define a one-time client action)
- DEFINE CLIENTOPT (Define an option to an option set)
- DEFINE CLOPTSET (Define a client option set name)
- DEFINE COLLOGGROUP (Define a collocation group)
- DEFINE COLLOCMEMBER (Define collocation group member)
- DEFINE COPYGROUP (Define a copy group)
- DEFINE DATAMOVER (Define a data mover)
- DEFINE DEVCLASS (Define a device class)
- DEFINE DOMAIN (Define a new policy domain)
- DEFINE DRIVE (Define a drive to a library)
- DEFINE EVENTSERVER (Define a server as the event server)
- DEFINE GRPMEMBER (Add a server to a server group)
- DEFINE LIBRARY (Define a library)
- DEFINE MACHINE (Define machine information for disaster recovery)
- DEFINE MACHNODEASSOCIATION (Associate a node with a machine)
- DEFINE MGMTCLASS (Define a management class)
- DEFINE NODEGROUP (Define a node group)
- DEFINE NODEGROUPMEMBER (Define node group member)
- DEFINE PATH (Define a path)
- DEFINE POLICYSET (Define a policy set)
- DEFINE PROFASSOCIATION (Define a profile association)
- DEFINE PROFILE (Define a profile)
- DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)
- DEFINE RECOVERYMEDIA (Define recovery media)
- DEFINE SCHEDULE (Define a client or an administrative command schedule)
- DEFINE SCRIPT (Define an IBM Spectrum Protect script)
- DEFINE SERVER (Define a server for server-to-server communications)
- DEFINE SERVERGROUP (Define a server group)
- DEFINE SPACETRIGGER (Define the space trigger)
- DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)
- DEFINE STGRULE (Define a rule for auditing storage pools)
- DEFINE STGRULE (Define a rule for generating data deduplication statistics)
- DEFINE STGRULE (Define a rule for reclaiming cloud containers)
- DEFINE STGRULE (Define a storage rule for tiering)
- DEFINE STGPOOL (Define a storage pool)
- DEFINE STGPOOLDIRECTORY (Define a storage pool directory)
- DEFINE SUBSCRIPTION (Define a profile subscription)
- DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)
- DEFINE VOLUME (Define a volume in a storage pool)

DEFINE ALERTTRIGGER (Define an alert trigger)

Use this command to trigger an alert whenever a server issues a specific error message. You can define a message number to be an alert trigger, assign it to a category, or specify administrators who can be notified of the alert by email.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-.,-----
      v          |
>>-Define ALERTTrigger-----message_number----->
      .-CAtegory--==--SErver-----
>--+-----+----->
      '-CAtegory--==--APplication--+'

```



```
define alerttrigger ANR1067E,ANR1073E Admin=BILL,DJADMIN
```

Related commands

Table 1. Commands related to DEFINE ALERTTRIGGER

Command	Description
DELETE ALERTTRIGGER (Remove a message from an alert trigger)	Removes a message number that can trigger an alert.
QUERY ALERTSTATUS (Query the status of an alert)	Displays information about alerts that have been issued on the server.
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	Displays message numbers that trigger an alert.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
UPDATE ALERTTRIGGER (Update a defined alert trigger)	Updates the attributes of one or more alert triggers.
UPDATE ALERTSTATUS (Update the status of an alert)	Updates the status of a reported alert.

DEFINE ASSOCIATION (Associate client nodes with a schedule)

Use this command to associate one or more clients with a schedule. You must assign a client node to the policy domain to which a schedule belongs. Client nodes process operations according to the schedules associated with the nodes.

Note:

1. IBM Spectrum Protect™ cannot run multiple schedules concurrently for the same client node.
2. In a macro, the server may stall if some commands (such as REGISTER NODE and DEFINE ASSOCIATION) are not committed as soon as you issue them. You could follow each command in a macro with a COMMIT command. However, a simpler solution is to include the -ITEMCOMMIT option with the DSMADMC command.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the policy domain to which the schedule belongs

Syntax

```
>>-DEFine ASSOCIation--domain_name--schedule_name----->
      .-,------
      v |
>----node_name+-----<<
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedule belongs.

schedule_name (Required)

Specifies the name of the schedule that you want to associate with one or more clients.

node_name (Required)

Specifies the name of a client node or a list of client nodes to associate with the specified schedule. Use commas to separate the items in the list. Do not leave spaces between the items and commas. You can use a wildcard character to specify a name. The command will not associate a listed client to the schedule if:

- The client is already associated with the specified schedule.
- The client is not assigned to the policy domain to which the schedule belongs.

- The client is a NAS node name. All NAS nodes are ignored.

Example: Associate client nodes with a schedule

Associate the client nodes SMITH or JOHN with the WEEKLY_BACKUP schedule. The associated clients are assigned to the EMPLOYEE_RECORDS policy domain.

```
define association employee_records
weekly_backup smith*,john*
```

Example: Associate client nodes with a schedule

Associate the client nodes JOE, TOM, and LARRY with the WINTER schedule. The associated clients are assigned to the EMPLOYEE_RECORDS policy domain; however, the client JOE is already associated with the WINTER schedule.

```
define association employee_records
winter joe,tom,larry
```

Related commands

Table 1. Commands related to DEFINE ASSOCIATION

Command	Description
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE ASSOCIATION	Deletes the association between clients and a schedule.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.
REGISTER NODE	Defines a client node to the server and sets options for that user.

DEFINE BACKUPSET (Define a backup set)

Use this command to define a client backup set that was previously generated on one server and make it available to the server that is running this command. The client node has the option of restoring the backup set from the server that is running this command rather than the one on which the backup set was generated.

Any backup set generated on one server can be defined to another server when the servers share a common device type. The level of the server to which the backup set is being defined must be equal to or greater than the level of the server that generated the backup set.

You can also use the DEFINE BACKUPSET command to redefine a backup set that was deleted on a server.

Privilege class

If the REQSYSAUTHOUTFILE server option is set to YES (the default), the administrator must have system privilege. If the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```

      .-,-----
      v'                               |
>>-DEFINE BACKUPSET-----+node_name-----+----->
                          '-node_group_name-'

>>-backup_set_name_prefix--DEVclass----device_class_name----->

      .-,-----
      v'                               |
>>-VOLumes-----volume_names----->
```


Specifies the number of days to retain the backup set on the server.

NOLimit

Specifies that the backup set must be retained on the server indefinitely.

If you specify NOLIMIT, IBM Spectrum Protect retains the volumes that contain the backup set forever, unless a user or administrator deletes the volumes from server storage.

DEscription

Specifies the description to associate with the backup set that belongs to the client node. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

WHERE DATAType

Specifies the backup sets containing the specified types of data are to be defined. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be defined. To specify multiple data types, separate the data types with commas and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be defined. ALL is the default value.

FILE

Specifies that a file level backup set is to be defined. File level backup sets contain files and directories that are backed up by the backup client.

IMAGE

Specifies that an image backup set is to be defined. Image backup sets contain images that are created by the backup-archive client BACKUP IMAGE command.

TOC

Specifies whether a table of contents (TOC) must be created for the file level backup set when it is defined. The TOC parameter is ignored when you define image and application data backup sets because a table of contents is always created for these backup sets.

Consider the following in determining whether you want to create a table of contents:

- If a table of contents is created, you can use the IBM Spectrum Protect web backup-archive client to examine the entire file system tree and choose files and directories to restore. Creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the management class that is specified by the TOCMGMTCLASS parameter. To create a table of contents extra processing, storage pool space, and possibly a mount point during the backup set operation is required.
- If a table of contents is not saved for a backup set, you can still restore individual files or directory trees by using the backup-archive client RESTORE BACKUPSET command if you know the fully qualified name of each file or directory to be restored.

This parameter is optional. The default value is Preferred. Possible values are:

No

Specifies that table of contents information is not saved for file level backup sets.

Preferred

Specifies that table of contents information must be saved for file level backup sets. However, a backup set does not fail just because an error occurs during creation of the table of contents.

Yes

Specifies that table of contents information must be saved for each file level backup set. A backup set fails if an error occurs during creation of the table of contents.

TOCMgmtclass

Specifies the name of the management class to which the table of contents must be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. In this case, creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the specified management class.

Example: Define a backup set

Define the PERS_DATA backup set that belongs to client node JANE to the server that is running this command. Retain the backup set on the server for 50 days. Specify that volumes VOL001 and VOL002 contain the data for the backup set. The volumes are to be read by a device that is assigned to the AGADM device class. Include a description.

```
define backupset jane pers_data devclass=agadm
volumes=vol1,vol2 retention=50
description="sector 7 base image"
```

Related commands

Table 1. Commands related to DEFINE BACKUPSET

Command	Description
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE NODEGROUP	Deletes a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DEFINE CLIENTACTION (Define a one-time client action)

Use this command to schedule one or more clients to process a command for a one-time action.

The server automatically defines a schedule and associates the client node to the schedule. The server assigns the schedule priority 1, sets the PERUNITS to ONETIME, and determines the number of days to keep the schedule active. The number of days is based on the value set with the SET CLIENTACTDURATION command.

How quickly the client processes this command depends on whether the scheduling mode for the client is set to server-prompted or client-polling. The client scheduler must be started on the client workstation in order for the server to process the schedule.

Remember: The start of the IBM Spectrum Protect™ scheduler depends on the processing of other threads in the server and other processes on the IBM Spectrum Protect server host system. The amount of time it takes to start the scheduler also depends on network traffic and how long it takes to open a socket, to connect with the IBM Spectrum Protect client, and to receive a response from the client. In general, the greater the processing and connectivity requirements on the IBM Spectrum Protect server and client, the longer it can take to start the scheduler.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy for the policy domain to which the schedule belongs.

Syntax

```

      .-,----- .
      v          |
>>- DEFine CLIENTAction-----node_name+----->

      .-Domain----*----- .
>-----+-----+----->
      |          .-,----- . |
      |          v          | |
      |'-Domain-----domain_name-+-'

      .-ACTion----Incremental----- .
```

```

>-----+----->
'-Action--==+Incremental-----+'
  +-Selective-----+
  +-Archive-----+
  |               | .-"-----.| |
  |               | '-SUBAction--==+-----+' |
  |               |               +-FASTBack----+ |
  |               |               +-SYSTEMState-+ |
  |               |               '-VM-----' |
  +-Backup-----+-----+
  |               | .-"-----.| |
  |               | '-SUBAction--==+-----+' |
  |               |               +-FASTBack----+ |
  |               |               +-SYSTEMState-+ |
  |               |               '-VM-----' |
  +-REStore-----+-----+
  +-RETRieve-----+-----+
  +-IMAGEBACKup-----+-----+
  +-IMAGERESTore-----+-----+
  +-Command-----+-----+
  '-Macro-----+'

>-----+----->
'-OPTions--===option_string-'

                               .-Wait----No-----.
>-----+-----+-----+----->>
'-OBJects--===object_string-' '-Wait--==+-No--+-'
                               '-Yes-'

```

Parameters

node_name (Required)

Specifies the name of the client node that will process the schedule associated with the action. If you specify multiple node names, separate the names with commas; do not use intervening spaces. You can use the asterisk wildcard character to specify multiple names.

DOmain

Specifies the list of policy domains used to limit the list of client nodes. Only client nodes that are assigned to one of the specified policy domains will be scheduled. All clients assigned to a matching domain will be scheduled. Separate multiple domain names with commas and no intervening spaces. If you do not specify a value, all policy domains will be included in the list.

ACTion

Specifies the action that occurs when this schedule is processed. Possible values are:

Incremental

Specifies that the schedule backs up all files that are new or that have changed since the last incremental backup. Incremental also backs up any file for which all existing backups might have expired.

Selective

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

Archive

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

Backup

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

REStore

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

RETRieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

Remember: A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.
IMAGERESTore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.
Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

SUBACTion

You can specify one of the following values:

""

When a null string (two double quotes) is specified with ACTION=BACKUP the backup is an incremental.

FASTBACK

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

SYSTEMState

Specifies that a client Systemstate backup is scheduled.

VApp

Specifies that a client vApp backup is scheduled. A vApp is a collection of pre-deployed virtual machines.

VM

Specifies that a client VMware backup operation is scheduled.

OPTions

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no effect when the client processes the scheduled command:

- MAXCMDRETRIES
- OPTFILE
- QUERYSCHEDPERIOD
- RETRYPERIOD
- SCHEDLOGNAME
- SCHEDMODE
- SERVERNAME
- TCPCLIENTADDRESS
- TCPCLIENTPORT

Windows When you define a scheduler service by using the DSMCUTIL command or the backup-archive client GUI wizard, you specify an options file. You cannot override the options in that options file by issuing the scheduled command. You must modify the options in your scheduler service.

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify `subdir=yes` and `domain all-local -systemobject`, enter:
 - `options='-subdir=yes -domain="all-local -c: -systemobject"'`
- To specify `domain all-local -c: -d:`, enter:
 - `options='-domain="all-local -c: -d:"'`

Windows Tip:

For Windows clients running in batch mode, if the use of quotation marks is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

OBJects

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when ACTION=INCREMENTAL. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify ACTION=INCREMENTAL without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:

```
ANS1082E Invalid destination file specification <filespec> entered.
```

- If you specify more than two file specifications, you receive this error:

```
ANS1102E Excessive number of command line arguments passed to the program!
```

When you specify ACTION=ARCHIVE, INCREMENTAL, or SELECTIVE for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

Windows If you are using characters that have a special meaning for Windows users, such as commas, surround the entire argument in two pairs of double quotes, then surround the entire string with single quotes. The following examples show you how to specify some file names:

- To specify C:\FILE 2, D:\GIF FILES, and E:\MY TEST FILE, enter:
 - OBJECTS="C:\FILE 2" "D:\GIF FILES" "E:\MY TEST FILE"
- To specify D:\TEST FILE, enter:
 - OBJECTS="'D:\TEST FILE'"
- To specify D:TEST,FILE:
 - OBJECTS="'"D:\TEST, FILE'"

AIX | **Linux** The following examples show how to specify some file names:

- To specify /home/file 2, /home/gif files, and /home/my test file, enter:
 - OBJECTS="'/home/file 2' '/home/gif files' '/home/my test file'"
- To specify /home/test file, enter:
 - OBJECTS="'/home/test file'"

Windows Tip:

For Windows clients running in batch mode, if the use of double quotes is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

Wait

Specifies whether to wait for a scheduled client operation to complete. This parameter is useful when defining client actions from a command script or macro. This parameter is optional. The default is No. Possible values are:

No

Specifies that you do not wait for the scheduled client operation to complete. If you specify this value and the value of the ACTION parameter is COMMAND, the return code indicates whether the client action was defined.

Yes

Specifies that you wait for the scheduled client operation to complete. If you specify this value and the value of the ACTION parameter is COMMAND, the return code indicates the status of the client operation.


```
define clientopt option_set_name inclexcl "include c:\proj\text\devel.*"
```

option_value (Required)

Specifies the value for the option. If the option includes more than one value, enclose the value in quotation marks.

Note:

1. The QUIET and VERBOSE options do not have an option value in the client option's file. To specify these values in a server client option set, specify a value of YES or NO.
2. To add an INCLUDE or EXCLUDE option for a file name that contains one or more spaces, put single quotation marks around the file specification, and double quotation marks around the entire option. See Example: Add an option to a client option set for more information.
3. The *option_value* is limited to 1024 characters.

Force

Specifies whether the server forces the client to use the option set value. The value is ignored for additive options, such as INCLEXCL and DOMAIN. The default is NO. This parameter is optional. The values are:

Yes

Specifies that the server forces the client to use the value. (The client cannot override the value.)

No

Specifies that the server does not force the client to use the value. (The client can override the value.)

SEQnumber

Specifies a sequence number when an option name is specified more than once. This parameter is optional.

Example: Add an option to a client option set

Add a client option (MAXCMDRETRIES 5) to a client option set named ENG.

```
define clientopt eng maxcmdretries 5
```

Example: Add an option to exclude a file from backup

Add a client option to the option set ENGBACKUP to exclude the c:\admin\file.txt from backup services.

```
define clientopt engbackup inclexcl "exclude c:\admin\file.txt"
```

Example: Add an option to exclude a directory from backup

Add a client option to the option set WINSPEC to exclude a temporary internet directory from backup services. When you use the EXCLUDE or INCLUDE option with file names that contain spaces, put single quotation marks around the file specification, then double quotation marks around the entire option.

```
define clientopt winspec inclexcl "exclude.dir '*:\...\Temporary Internet Files'"
```

Example: Add an option to bind files in specified directories

Add client options to the option set WINSPEC to bind all files in directories C:\Data and C:\Program Files\My Apps to a management class named PRODCLASS.

```
define clientopt winspec inclexcl "include C:\Data\...\* prodclass"  
define clientopt winspec inclexcl "include 'C:\Program  
Files\My Apps\...\*' prodclass"
```

Related commands

Table 1. Commands related to DEFINE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.

Command	Description
REGISTER NODE	Defines a client node to the server and sets options for that user.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.
UPDATE NODE	Changes the attributes that are associated with a client node.

DEFINE CLOPTSET (Define a client option set name)

Use this command to define a name for a set of options you can assign to clients for archive, backup, restore, and retrieve operations.

To add options to the new set, issue the DEFINE CLIENTOPT command.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax

```
>>-DEFine CLOptset--option_set_name----->
>--+-----+-----><
  '-DESCRiption----description-'
```

Parameters

option_set_name (Required)

Specifies the name of the client option set. The maximum length of the name is 64 characters.

DESCRiption

Specifies a description of the client option set. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. This parameter is optional.

Example: Define a client option set

To define a client option set named ENG issue the following command.

```
define cloptset eng
```

Related commands

Table 1. Commands related to DEFINE CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DEFINE COLLOGROUP (Define a collocation group)

Use this command to define a collocation group. A *collocation group* is a group of nodes or file spaces on a node whose data is collocated on a minimal number of sequential access volumes. Their data is collocated only if the storage pool definition is set to collocate by group (COLLOCATE=GROUP).

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
>>-DEfIne COLLOGGroup--group_name----->
>--+-----+-----><
  '-DEScRiption----description-'
```

Parameters

group_name

Specifies the name of the collocation group name that you want to create. The maximum length of the name is 30 characters.

DEScRiption

Specifies a description of the collocation group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Define a collocation group

To define a node or file space collocation group named GROUP1, issue the following command:

```
define collogroup group1
```

Related commands

Table 1. Commands related to DEFINE COLLOGROUP

Command	Description
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DEFINE COLLOCMEMBER (Define collocation group member)

Issue this command to add a client node to a collocation group or to add a file space from a node to a collocation group. A collocation group is a group of nodes or file spaces on a node whose data is collocated on a minimal number of sequential access volumes.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

Add a node to a collocation group

```
          .-,------.
          v            |
>>-DEfIne COLLOCMember--group_name----node_name-+-----><
```

Parameters

group_name

Specifies the name of the collocation group to which you want to add a client node.

node_name

Specifies the name of the client node that you want to add to the collocation group. You can specify one or more names. Separate multiple names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple names.

Add a file space from a node to a collocation group

```
>>-DEfIne COLLOCMember--group_name--node_name----->
```

```
          .-,------.
          v            |
>>-Filespace-----file_space_name-+----->

.-NAMEType-----SERVER-----.
>--+-----+----->
'-NAMEType-----+SERVER--+-'
      +-UNICODE-+
      '-FSID----'
```

```
.-CODEType-----BOTH-----.
>--+-----+-----><
'-CODEType-----+BOTH-----+'
      +-UNICODE-----+
      '-NONUNICODE-'
```

Parameters

group_name

Specifies the name of the collocation group to which you want to add a file space.

node_name

Specifies the client node where the file space is located.

Filespace

Specifies the *file_space_name* on the client node that you want to add to the collocation group. You can specify one or more file space names that are on a specific client node. If you specify multiple file space names, separate the names with commas with no intervening spaces. You can also use wildcard characters to specify multiple file space names. For example:

```
define collocmember manufacturing linux237 filespace=*_linux_fs
```

This command places all file spaces on the linux237 node with a name that ends with `_linux_fs` into the manufacturing collocation group.

See the following list for tips about working with collocation groups:

- When you add members to a new collocation group, the type of the first collocation group member determines the type of the collocation group. The group can either be a node collocation group or a file space collocation group. Restriction: After the collocation group type is set, it cannot be changed.
- You cannot mix collocation group member types when you add members to a collocation group (either a node group or a file space group).
- For a file space collocation group, you can add file spaces to the group. The file spaces must use the same value as the `node_name` parameter that is specified when the collocation group is established.
- A client node can be included in multiple file space groups. However, if a node is a member of a node collocation group, it cannot be a member of a file space collocation group.
- A file space can be a member of only one file space group.

NAMETYPE

Specify how you want the server to interpret the file space names that you enter. Specify this parameter when the server communicates with clients that have Unicode support. A backup-archive client with Unicode support is available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare systems. The filespace name cannot be a wildcard character when NAMETYPE is specified for a filespace collocation group. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. Whether the name can be converted depends on the characters in the names and the server code page. Conversion might fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names by their file space IDs (FSIDs).

CODETYPE

Specify how you want the server to interpret the file space names that you enter. Use this parameter when you use a wildcard character for the file space name. For example:

```
define collocmember production Win_3419 filespace=* codetype=unicode
```

This example command adds all file spaces from the Win_3419 node to the production collocation group. The default is BOTH, so the file spaces are included, regardless of code page type. You can specify one of the following values:

BOTH

Include the file spaces, regardless of code page type.

UNICODE

Include file spaces that are only in Unicode.

NONUNICODE

Include file spaces that are not in Unicode.

Define two collocation group members

Define two members, NODE1 and NODE2, to a collocation group, GROUP1.

```
define collocmember group1 node1,node2
```

Define one file space group member CNTR90524, on node clifton to collocation group TSM_alpha_1

```
define collocmember TSM_alpha_1 clifton filespace=CNTR90524
```

Related commands

Table 1. Commands related to DEFINE COLLOCMEMBER

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DEFINE COPYGROUP (Define a copy group)

Use this command to define a new backup or archive copy group within a specific management class, policy set, and policy domain. The server uses the backup and archive copy groups to control how clients back up and archive files, and to manage the backed-up and archived files.

To enable clients to use the new copy group, you must activate the policy set that contains the new copy group.

You can define one backup and one archive copy group for each management class. To ensure that client nodes can back up files, include a backup copy group in the default management class for a policy set.

Attention: The DEFINE COPYGROUP command fails if you specify a copy storage pool as a destination.

The DEFINE COPYGROUP command has two forms, one for defining a backup copy group and one for defining an archive copy group. The syntax and parameters for each form are defined separately.

Table 1. Commands related to DEFINE COPYGROUP

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
BACKUP NODE	Backs up a network-attached storage (NAS) node.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.

Command	Description
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
SET ARCHIVERETENTIONPROTECTION	Specifies whether data retention protection is activated.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

- DEFINE COPYGROUP (Define a backup copy group)
Use this command to define a new backup copy group within a specific management class, policy set, and policy domain.
- DEFINE COPYGROUP (Define an archive copy group)
Use this command to define a new archive copy group within a specific management class, policy set, and policy domain.

DEFINE COPYGROUP (Define a backup copy group)

Use this command to define a new backup copy group within a specific management class, policy set, and policy domain.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax

```
>>-DEFine COpYgroup--domain_name--policy_set_name--class_name--->
    .-STANDARD-.  .-Type----Backup-.
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-STANDARD-'  '-Type----Backup-'
                                     .-FREQuency----0----.
>--DESTination---pool_name---+-----+-----+-----+----->
                                     '-FREQuency----days-'
    .-VERExists----2-----+-----+-----+-----+----->
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-VERExists---+number--+-'
                                     '-NOLimit-'
    .-VERDeleted----1-----+-----+-----+-----+----->
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-VERDeleted---+number--+-'
                                     '-NOLimit-'
    .-RETEExtra---30-----+-----+-----+-----+----->
    .-RETOOnly---60-----+-----+-----+-----+----->
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-RETEExtra---+days---+-'  '-RETOOnly---+days---+-'
                                     '-NOLimit-'  '-NOLimit-'
    .-MODE----MODified----+-----+-----+-----+----->
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-MODE---+MODified--+-'
                                     '-ABSolute-'
    .-SERialization---SHRStatic-----+-----+-----+----->
>--+-----+-----+-----+-----+-----+-----+-----+----->
    '-SERialization---+SHRStatic--+-'
                                     +-STatic-----+
                                     +-SHRDYnamic+
                                     '-DYnamic----'
>--+-----+-----+-----+-----+-----+-----+-----+-----><
```

'-TOCDestination--=-----pool_name---'

Parameters

domain_name (Required)

Specifies the policy domain for which you are defining the copy group.

policy_set_name (Required)

Specifies the policy set for which you are defining the copy group.

You cannot define a copy group for a management class that belongs to the ACTIVE policy set.

class_name (Required)

Specifies the management class for which you are defining the copy group.

STANDARD

Specifies the name of the copy group, which must be STANDARD. This parameter is optional. The default value is STANDARD.

Type=Backup

Specifies that you want to define a backup copy group. The default parameter is BACKUP. This parameter is optional.

DESTINATION (Required)

Specifies the primary storage pool where the server initially stores backup data. You cannot specify a copy storage pool as the destination.

FREQUENCY

Specifies how frequently IBM Spectrum Protect™ can back up a file. This parameter is optional. IBM Spectrum Protect backs up a file only when the specified number of days has elapsed since the last backup. The FREQUENCY value is used only during a full incremental backup operation. This value is ignored during selective backup or partial incremental backup. You can specify an integer from 0 to 9999. The default value is 0, meaning that IBM Spectrum Protect can back up a file regardless of when the file was last backed up.

VERExists

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional. The default value is 2.

If an incremental backup operation causes the limit to be exceeded, the server expires the oldest backup version that exists in server storage. Possible values are:

number

Specifies the number of backup versions to retain for files that are currently on the client file system. You can specify an integer from 1 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 2. Preferred values are 3, 4, or more.

NOLimit

Specifies that you want the server to retain all backup versions.

The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using IBM Spectrum Protect. This parameter is optional. The default value is 1.

If a user deletes a file from the client file system, the next incremental backup causes the server to expire the oldest versions of the file in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter. Possible values are:

number

Specifies the number of backup versions to retain for files that are deleted from the client file system after being backed up. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEExtra

Specifies the number of days to retain a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes inactive versions based on retention time even if the number of

inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. The default value is 30 days. Possible values are:

days

Specifies the number of days to retain inactive backup versions. You can specify an integer from 0 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 14 days. The preferred value is 30 or more days.

NOLimit

Specifies that you want to retain inactive backup versions indefinitely.

If you specify NOLIMIT, the server deletes inactive backup versions based on the VEREXISTS parameter (when the file still exists on the client file system) VERDELETED parameter (when the file no longer exists on the client file system).

REOnly

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. The default value is 60. Possible values are:

days

Specifies the number of days to retain the last remaining inactive version of a file. You can specify an integer from 0 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 30 days.

NOLimit

Specifies that you want to keep the last remaining inactive version of a file indefinitely.

If you specify NOLIMIT, the server retains the last remaining backup version forever, unless a user or administrator deletes the file from server storage.

MODE

Specifies whether IBM Spectrum Protect backs up a file only if the file has changed since the last backup, or whenever a client requests a backup. This parameter is optional. The default value is MODIFIED. Possible values are:

MODified

Specifies that IBM Spectrum Protect backs up the file only if it has changed since the last backup. IBM Spectrum Protect considers a file changed if any of the following is true:

- The date last modified is different
- The file size is different
- The file owner is different
- The file permissions are different

ABSolute

Specifies that IBM Spectrum Protect backs up the file regardless of whether it has been modified.

The MODE value is used only for full incremental backup. This value is ignored during partial incremental backup or selective backup.

SERialization

Specifies how IBM Spectrum Protect processes files or directories when they are modified during backup processing. This parameter is optional. The default value is SHRSTATIC. Possible values are:

SHRStatic

Specifies that IBM Spectrum Protect backs up a file or directory only if it is not being modified during backup. IBM Spectrum Protect attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file or directory is modified during each backup attempt, IBM Spectrum Protect does not back it up.

Static

Specifies that IBM Spectrum Protect backs up a file or directory only if it is not being modified during backup. IBM Spectrum Protect attempts to perform the backup only once.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file or directory is being modified during a backup attempt, IBM Spectrum Protect backs up the file or directory during the last attempt even though the file or directory is being modified. IBM Spectrum Protect attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that IBM Spectrum Protect backs up a file or directory on the first attempt, regardless of whether the file or directory is being modified during backup processing.

Attention: Be careful about using the SHRDYNAMIC and DYNAMIC values. IBM Spectrum Protect uses these values to determine if it backs up a file or directory while modifications are occurring. As a result, the backup version might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file or directory because it contains some, but not all, modifications. If a file that contains a fuzzy backup is restored, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that IBM Spectrum Protect creates a backup version only if the file or directory is not being modified.

TOCDestination

Specifies the primary storage pool in which a table of contents (TOC) will initially be stored for any Network Data Management Protocol (NDMP) backup or backup set operation for which a TOC is generated. This parameter is optional. You cannot specify a copy storage pool as the destination. The storage pool specified for the destination must have NATIVE or NONBLOCK data format. To avoid mount delays, it is recommended that the storage pool have a device class of DISK or DEVTYPE=FILE. TOC generation is an option for NDMP backup operations, but is not supported for other image-backup operations.

If TOC creation is requested for a backup operation that uses NDMP and the image is bound to a management class whose backup copy group does not specify a TOC destination, the outcome will depend on the TOC parameter for the backup operation.

- If TOC=PREFERRED (the default), the backup proceeds without creation of a TOC.
- If TOC=YES, the entire backup fails because no TOC can be created.

Example: Create a backup copy group

Create a backup copy group named STANDARD for management class ACTIVEFILES in policy set VACATION in the EMPLOYEE_RECORDS policy domain. Set the backup destination to BACKUPPOOL. Set the minimum interval between backups to three days, regardless of whether the files have been modified. Retain up to five backup versions of a file while the file exists on the client file system.

```
define copygroup employee_records
vacation activefiles standard type=backup
destination=backuppools frequency=3
verexists=5 mode=absolute
```

DEFINE COPYGROUP (Define an archive copy group)

Use this command to define a new archive copy group within a specific management class, policy set, and policy domain.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax

```
>>-DEFine COpYgroup--domain_name--policy_set_name--class_name-->
. -STANDARD- .
>--+-----+---Type-----Archive--DESTination-----pool_name----->
' -STANDARD- '

. -FREquency-----Cmd- . . -RETVer-----365----- .
>--+-----+-----+-----+-----+-----+-----+----->
' -FREquency-----Cmd- ' ' -RETVer-----+--days-----+ '
' -NOLimit- '
```

```

.-REtInit---CREAtion-. .-REtMin---365-----.
>-----+-----+-----+----->
'-REtInit---EvEnt---' '-REtMin---days---'

.-MODE---ABSolute-.
>-----+-----+-----+----->
'-MODE---ABSolute-'

.-SERialization---SHRStatic-----.
>-----+-----+-----+-----><
'-SERialization---SHRStatic---+'
                    +-Static-----+
                    +-SHRDYnamic-+-
                    '-DYnamic-----'

```

Parameters

domain_name (Required)

Specifies the name of the policy domain for which you are defining the copy group.

policy_set_name (Required)

Specifies the name of the policy set for which you are defining the copy group.

You cannot define a copy group for a management class that belongs to the ACTIVE policy set.

class_name (Required)

Specifies the name of the management class for which you are defining the copy group.

STANDARD

Specifies the name of the copy group, which must be STANDARD. This parameter is optional. The default value is STANDARD.

Type=Archive (Required)

Specifies that you want to define an archive copy group.

DESTination (Required)

Specifies the primary storage pool where the server initially stores the archive copy. You cannot specify a copy storage pool as the destination.

FREQuency=Cmd

Specifies the copy frequency, which must be CMD. This parameter is optional. The default value is CMD.

REtVer

Specifies the number of days to keep an archive copy. This parameter is optional. The default value is 365. Possible values are:

days

Specifies the length of time to keep an archive copy. You can specify an integer in the range 0 - 30000.

Tip: To help ensure that your data can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 30 days.

The RETENTIONEXTENSION server option can affect the volume retention if the following conditions are true:

- You specify zero for the number of days
- The destination storage pool for the archive copy group is a SnapLock storage pool (RECLAMATIONTYPE=SNAPLOCK)

If the two conditions are met, retention of the volumes is defined by the value of the RETENTIONEXTENSION server option. The RETENTIONEXTENSION server option value also applies if data is copied or moved into the SnapLock storage pool by a server process such as migration, or by using the MOVE DATA or MOVE NODEDATA commands.

NOLimit

Specifies that you want to keep an archive copy indefinitely.

If you specify NOLIMIT, the server retains archive copies forever, unless a user or administrator deletes the file from server storage. If you specify NOLIMIT, you cannot also specify EVENT for the RETINIT parameter.

The value of the RETVER parameter can affect the management class to which the server binds an archived directory. If the client does not use the ARCHMC option, the server binds directories that are archived to the default management class. If the default management class has no archive copy group, the server binds directories that are archived to the management class with the shortest retention period.

The RETVER parameter of the archive copy group of the management class to which an object is bound determines the retention criterion for each object. See the SET ARCHIVERETENTIONPROTECTION command for a description of data protection.

If the primary storage pool specified in the DESTINATION parameter belongs to a Centera device class and data protection is enabled, then the RETVER value is sent to Centera for retention management purposes. See the SET ARCHIVERETENTIONPROTECTION command for a description of data protection.

RETInit

Specifies when the retention time specified by the RETVER attribute is initiated. This parameter is optional. If you define the RETINIT value during copy group creation, you cannot modify it later. The default value is CREATION. Possible values are:

CREATion

Specifies that the retention time specified by the RETVER attribute is initiated at the time an archive copy is stored on the IBM Spectrum Protect™ server.

EVent

Specifies that the retention time specified in the RETVER parameter is initiated at the time a client application notifies the server of a retention-initiating event for the archive copy. If you specify RETINIT=EVENT, you cannot also specify RETVER=NOLIMIT.

Tip: You can place a deletion hold on an object that was stored with RETINIT=EVENT for which the event has not been signaled. If the event is signaled while the deletion hold is in effect, the retention period is initiated, but the object is not deleted while the hold is in effect.

RETMIn

Specifies the minimum number of days to keep an archive copy after it is archived. This parameter is optional. The default value is 365. If you specify RETINIT=CREATION, this parameter is ignored.

MODE=ABSolute

Specifies that a file is always archived when the client requests it. The MODE must be ABSOLUTE. This parameter is optional. The default value is ABSOLUTE.

SERialization

Specifies how IBM Spectrum Protect processes files that are modified during archive. This parameter is optional. The default value is SHRSTATIC. Possible values are:

SHRStatic

Specifies that IBM Spectrum Protect archives a file only if it is not being modified. IBM Spectrum Protect attempts to perform an archive operation as many as four times, depending on the value that is specified for the CHANGINGRETRIES client option. If the file is modified during the archive attempt, IBM Spectrum Protect does not archive the file.

Static

Specifies that IBM Spectrum Protect archives a file only if it is not being modified. IBM Spectrum Protect attempts to perform the archive operation only once.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDYnamic

Specifies that if the file is being modified during an archive attempt, IBM Spectrum Protect archives the file during its last attempt even though the file is being modified. IBM Spectrum Protect attempts to archive the file as many as four times, depending on the value that is specified for the CHANGINGRETRIES client option.

DYnamic

Specifies that IBM Spectrum Protect archives a file on the first attempt, regardless of whether the file is being modified during archive processing.

Attention: Be careful about using the SHRDYNAMIC and DYNAMIC values. IBM Spectrum Protect uses them to determine if it archives a file while modifications are occurring. As a result, the archive copy might be a fuzzy backup. A fuzzy backup does not accurately reflect what is in the file because it contains some, but not all, modifications. If a file that contains a fuzzy backup is retrieved, the file might or might not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that IBM Spectrum Protect creates an archive copy only if the file is not being modified.

Example: Define an archive copy group for event-based retention

Create an archive copy group named STANDARD for management class EVENTMC in policy set SUMMER in the PROG1 policy domain. Set the archive destination to ARCHIVEPOOL, where the archive copy is kept until the server is notified of an event to

initiate the retention time, after which the archive copy is kept for 30 days. The archive copy will be kept for a minimum of 90 days after being stored on the server, regardless of when the server is notified of an event to initiate the retention time.

```
define copygroup prog1 summer eventmc standard type=archive
destination=archivepool retinit=event retver=30 retmin=90
```

DEFINE DATAMOVER (Define a data mover)

Use this command to define a data mover. A data mover is a named device that accepts a request from IBM Spectrum Protect™ to transfer data. A data mover can be used to complete outboard copy operations.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DATAMover--data_mover_name----->
. -Type-----NAS-----
>--+-----+-----HLAddress-----address-->
|                                     (1) (2) |
' -Type-----+--NASCLUSTER--+-----'
      '-NASVSERVER-'

. -LLAddress-----10000----.
>--+-----+-----USERid-----userid----->
' -LLAddress-----tcp_port-'

. -ONLine-----Yes-----.
>--PASsword-----password--+-----+----->
      '-ONLine-----+--Yes+-'
      '-No--'

>--DATAFormat-----+--NETAPPDump--+-----><
      +-CELERRADump-+
      '-NDMPDump-----'
```

Notes:

1. You can specify `TYPE=NASCLUSTER` and `TYPE=NASVSERVER` only on an AIX®, Linux, or Windows operating system.
2. You can specify `TYPE=NASCLUSTER` and `TYPE=NASVSERVER` only if `DATAFORMAT=NETAPPDUMP`.

Parameters

`data_mover_name` (Required)

Specifies the name of the data mover. This name must be the same as a node name that you previously registered by using the `REGISTER NODE TYPE=NAS` command. The data that is backed up from this NAS data mover will be assigned to this node name in the server database. A maximum of 64 characters can be used to specify the name.

Type

Specifies the type of data mover. This parameter is optional. The default value is `NAS`.

`NAS`

Specifies that the data mover is a NAS file server.

`NASCLUSTER`

Specifies that the data mover is a clustered NAS file server.

Restriction: You can specify the `NASCLUSTER` value only if `DATAFORMAT=NETAPPDUMP`.

`NASVSERVER`

Specifies that the data mover is a virtual storage device within a cluster.

Restriction: You can specify the `NASVSERVER` value only if `DATAFORMAT=NETAPPDUMP`.

`HLAddress` (Required)

Specifies either the numerical IP address or the domain name that is used to access the NAS file server.

Tip: To determine the numerical IP address, access the NAS file server. Then, follow the instructions in the file server documentation for obtaining the address.

LLAddress

Specifies the TCP port number to access the NAS device for Network Data Management Protocol (NDMP) sessions. This parameter is optional. The default value is 10000.

USERid (Required)

Specifies the user ID for a user that is authorized to initiate an NDMP session with the NAS file server. For example, enter the user ID that is configured on the NetApp file server for NDMP connections.

Tip: To determine the user ID, access the NAS file server. Then, follow the instructions in the file server documentation for obtaining the user ID.

PASsword (Required)

Specifies the password for the user ID to log on to the NAS file server.

Tip: To determine the password, access the NAS file server. Then, follow the instructions in the file server documentation for obtaining the password.

ONLine

Specifies whether the data mover is available for use. This parameter is optional. The default is YES.

Yes

The default value. Specifies that the data mover is available for use.

No

Specifies that the data mover is not available for use. When the hardware is being maintained, you can use the UPDATE DATAMOVER command to set the data mover offline.

If a library is controlled by using a path from a NAS data mover to the library, and the NAS data mover is offline, the server is not able to access the library. If the server is halted and restarted while the NAS data mover is offline, the library is not initialized.

DATAFormat (Required)

Specifies the data format that is used by this data mover.

NETAPPDump

Must be used for NetApp NAS file servers and the IBM® System Storage® N Series.

CELERRADump

Must be used for EMC Celerra NAS file servers.

NDMPDump

Must be used for NAS file servers other than NetApp or EMC file servers.

Example: Define a data mover by domain name

Define a data mover for the node named NAS1. The domain name for the data mover is NETAPP2.EXAMPLE.COM at port 10000.

```
define datamover nas1 type=nas hladdress=netapp2.example.com lladdress=10000
userid=root password=admin dataformat=netappdump
```

Example: Define a data mover by IP address

Define a data mover for the node named NAS2. The numerical IP address for the data mover is 203.0.113.0, at port 10000. The NAS file server is not a NetApp or EMC file server.

```
define datamover nas2 type=nas hladdress=203.0.113.0 lladdress=10000
userid=root password=admin dataformat=ndmpdump
```

Example: Define a data mover for a clustered file server by IP address

Define a data mover for the clustered file server named NAS3. The NAS file server is a NetApp device. The numerical IP address for the data mover is 198.51.100.0, at port 10000.

```
define datamover nas3 type=nascluster hladdress=198.51.100.0
lladdress=10000 userid=root password=admin dataformat=netappdump
```

Related commands

Table 1. Commands related to DEFINE DATAMOVER

Command	Description
DEFINE PATH	Defines a path from a source to a destination.
DELETE DATAMOVER	Deletes a data mover.
QUERY DATAMOVER	Displays data mover definitions.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE DATAMOVER	Changes the definition for a data mover.

DEFINE DEVCLASS (Define a device class)

Use this command to define a device class for a type of storage device. The server requires that a device class be defined to allow the use of a device.

For the most up-to-date list of supported devices and valid device class formats, see the IBM Spectrum Protect™ Supported Devices website: [AIX](#) | [Windows](#)

- Supported devices for AIX and Windows

Linux

- Supported devices for Linux

Note: The DISK device class is defined by IBM Spectrum Protect and cannot be modified with the DEFINE DEVCLASS command.

[AIX](#) | [Linux](#) If you are defining a device class for devices that are to be accessed through a z/OS® media server, see Define device class for z/OS media server.

The following IBM Spectrum Protect device classes are ordered by device type.

- r_cmd_devclass_3590_define.dita#r_cmd_devclass_3590_define
- r_cmd_devclass_3592_define.dita#r_cmd_devclass_3592_define
- r_cmd_devclass_4mm_define.dita#r_cmd_devclass_4mm_define
- r_cmd_devclass_8mm_define.dita#r_cmd_devclass_8mm_define
- r_cmd_devclass_centera_define.dita#r_cmd_devclass_centera_define
- r_cmd_devclass_dlt_define.dita#r_cmd_devclass_dlt_define
- r_cmd_devclass_ecartridge_define.dita#r_cmd_devclass_ecartridge_define
- r_cmd_devclass_file_define.dita#r_cmd_devclass_file_define
- [AIX](#) | [Windows](#) r_cmd_devclass_generictape_define.dita#r_cmd_devclass_generictape_define
- r_cmd_devclass_lto_define.dita#r_cmd_devclass_lto_define
- r_cmd_devclass_nas_define.dita#r_cmd_devclass_nas_define
- r_cmd_devclass_removablefile_define.dita#r_cmd_devclass_removablefile_define
- r_cmd_devclass_server_define.dita#r_cmd_devclass_server_define
- r_cmd_devclass_volsafe_define.dita#r_cmd_devclass_volsafe_define

Table 1. Commands related to DEFINE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up IBM Spectrum Protect device information to a file.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE DEVCLASS	Changes the attributes of a device class.

DEFINE DEVCLASS (Define a 3590 device class)

Use the 3590 device class when you are using 3590 tape devices.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see DEFINE DEVCLASS (Define a 3590 device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRary----library_name--DEVType-----3590----->
. -FORMAT----DRIVE-----.
>--+-----+-----+-----+-----+----->
' -FORMAT----+DRIVE---+' ' -ESTCAPacity---size-'
      +-3590B---+
      +-3590C---+
      +-3590E-B-+
      +-3590E-C-+
      +-3590H-B-+
      '-3590H-C-'

. -PREFIX----ADSM-----.
>--+-----+-----+-----+-----+----->
' -PREFIX----+ADSM-----+'
      '-tape_volume_prefix-'

. -MOUNTRetention---60-----. -MOUNTWait---60-----.
>--+-----+-----+-----+-----+----->
' -MOUNTRetention---minutes-' ' -MOUNTWait---minutes-'

. -MOUNTLimit----DRIVES-----.
>--+-----+-----+-----+-----+-----><
' -MOUNTLimit----+DRIVES-+-'
      +-number-+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRary (Required)

Specifies the name of the defined library object that contains the tape drives that can be used by this device class.

For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=3590 (Required)

Specifies the 3590 device type is assigned to the device class. 3590 indicates that IBM® 3590 cartridge tape devices are assigned to this device class.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other

must have LTO-6 drives and media.

The following tables list the recording formats, estimated capacities, and recording format options for 3590 devices:

Table 1. Recording formats and default estimated capacities for 3590

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
3590B	10.0 GB	Uncompressed (basic) format
3590C	See note 20.0 GB	Compressed format
3590E-B	10.0 GB	Uncompressed (basic) format, similar to the 3590B format
3590E-C	See note 20.0 GB	Compressed format, similar to the 3590C format
3590H-B	30.0 GB (J cartridge – standard– length) 60.0 GB (K cartridge - extended length)	Uncompressed (basic) format, similar to the 3590B format
3590H-C	See note 60.0 GB (J cartridge - standard length) 120.0 GB (K cartridge - extended length)	Compressed format, similar to the 3590C format

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be greater than the listed value.

Table 2. 3590 device recording format selections

Device	Format					
	3590B	3590C	3590E-B	3590E-C	3590H-B	3590H-C
3590	Read/Write	Read/Write	–	–	–	–
Ultra SCSI	Read/Write	Read/Write	–	–	–	–
3590E	Read	Read	Read/Write	Read/Write	–	–
3590H	Read	Read	Read	Read	Read/Write	Read/Write

ESTCAPACITY

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADMS. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is AD\$M.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DEFINE DEVCLASS (Define a 3592 device class)

Use the 3592 device class when you are using 3592 tape devices.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see DEFINE DEVCLASS (Define a 3592 device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRary-----library_name--DEVType-----3592----->

                                     (1)
.-LBProtect-----No----- .-WORM-----No-----
>--+-----+-----+-----+----->
'-LBProtect-----+READWrite+-' '-WORM-----+Yes+-'
      +-WRITEOnly+
      '-No-----'
                                     '-No--'

.-SCALECAPacity-----100----- .-FORMAT-----DRIVE-----
>--+-----+-----+-----+----->
'-SCALECAPacity-----+100+-' '-FORMAT-----+DRIVE-----+'
      +-90--+
      '-20--'
                                     +-3592-----+
                                     +-3592C----+
                                     +-3592-2----+
                                     +-3592-2C--+
                                     +-3592-3----+
                                     +-3592-3C--+
                                     +-3592-4----+
                                     +-3592-4C--+
                                     +-3592-5----+
                                     +-3592-5C--+
                                     +-3592-5A--+
                                     '-3592-5AC-'

>--+-----+-----+-----+----->
'-ESTCAPacity-----size-'

.-PREFIX-----ADSM-----
>--+-----+-----+-----+----->
'-PREFIX-----+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTRetention-----60----- .-MOUNTWait-----60-----
>--+-----+-----+-----+----->
'-MOUNTRetention-----minutes-' '-MOUNTWait-----minutes-'

.-MOUNTLimit-----DRIVES-----
>--+-----+-----+-----+----->
'-MOUNTLimit-----+DRIVES+-'
      +-number+
      '-0-----'

                                     (1) (2)
.-DRIVEEncryption-----ALLOW-----
>--+-----+-----+-----+-----><
'-DRIVEEncryption-----+ON-----+'
      +-ALLOW-----+
      +-EXTERNAL+
      '-OFF-----'
```

Notes:

1. You cannot specify both WORM=Yes and DRIVEENCRYPTION=ON.
2. Drive encryption is supported only for 3592 Generation 2 or later drives.

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRary (Required)

Specifies the name of the defined library object that contains the tape drives that can be used by this device class.

For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=3592 (Required)

Specifies that the 3592 device type is assigned to the device class.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When LBPROTECT is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is NO.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect™ and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the BACKUP DB command.

When the LBPROTECT parameter is set to READWRITE, you do not have to specify the CRCDATA parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the BACKUP DB command.

No

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on IBM® 3592 Generation 3 drives and later with 3592 Generation 2 media and later.

See Technote 1634851, Additional information on the IBM Spectrum Protect LBProtect option, for an explanation about when to use the LBProtect parameter.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is No. The field can contain one of the following values:

Yes

Specifies that the drives use WORM media.

No

Specifies that the drives do not use WORM media.

Remember:

1. To use 3592 WORM support in 3584 libraries, you must specify the WORM parameter. The server distinguishes between WORM and non-WORM scratch volumes. However, to use 3592 WORM support in 349X libraries, you also must set the WORMSCRATCHCATEGORY on the DEFINE LIBRARY command. For details, see DEFINE LIBRARY (Define a library).
2. When WORM=Yes, the only valid value for the SCALECAPACITY parameter is 100.
3. Verify with your hardware vendors that your hardware is at the appropriate level of support.

SCALECAPacity

Specifies the percentage of the media capacity that can be used to store data. This parameter is optional. The default is 100. Possible values are 20, 90, or 100.

Setting the scale capacity percentage to 100 provides maximum storage capacity. Setting it to 20 provides fastest access time.

Note: The scale capacity value takes effect only when data is first written to a volume. Any updates to the device class for scale capacity do not affect volumes that already have data that is written to them until the volume is returned to scratch status.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats, estimated capacities, and recording format options for 3592 devices.

Tip: The format name is specified as, for example, 3592-X, 3592-XC, 3592-XA, or 3592-XAC, where X indicates the drive generation, C indicates a compressed format, and A indicates an archive drive.

Table 1. Recording formats and default estimated capacities for 3592

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
3592	300 GB	Uncompressed (basic) format
3592C	See note.	Compressed format
3592-2	500 GB	Uncompressed (basic) format JA tapes
	700 GB	Uncompressed (basic) format JB tapes
3592-2C	1.5 TB	Compressed format JA tapes
	2.1 TB	Compressed format JB tapes
3592-3	640 GB	Uncompressed (basic) format JA tapes
	1 TB	Uncompressed (basic) format JB tapes
3592-3C	1.9 TB	Compressed format JA tapes
	3 TB	Compressed format JB tapes
3592-4	400 GB	Uncompressed (basic) format JK tapes
	1.5 TB	Uncompressed (basic) format JB tapes
	3.1 TB	Uncompressed (basic) format JC tapes
3592-4C	1.2 TB	Compressed format JK tapes
	4.4 TB	Compressed format JB tapes
	9.4 TB	Compressed format JC tapes

Format	Estimated capacity	Description
3592-5	900 GB	Uncompressed (basic) format JK tapes
(For IBM TS1150 Model 3592 E08 drives with product ID 03592E08)	7 TB	Uncompressed (basic) format JC/JY tapes
	2 TB	Uncompressed (basic) format JL tapes
	10 TB	Uncompressed (basic) format JD/JZ tapes
3592-5C	Depends on the compressibility of the data	Compressed format JK tapes
(For IBM TS1150 Model 3592 E08 drives with product ID 03592E08)		Compressed format JC/JY tapes
		Compressed format JL tapes
		Compressed format JD/JZ tapes
3592-5A	3 TB	Uncompressed (basic) format JL tapes
(For IBM TS1155 Model 3592 55F drives with product ID 0359255F)	15 TB	Uncompressed (basic) format JD/JZ tapes
	3592-5AC	Depends on the compressibility of the data
(For IBM TS1155 Model 3592 55F drives with product ID 0359255F)	Compressed format JD/JZ tapes	

Note: If this format uses the compression feature for tape drives, depending on the effectiveness of compression, the actual capacity might be different from the estimated capacity.

Important: For optimal performance, avoid mixing different generations of drives in a single SCSI library. If you must mix drive generations in a SCSI library, use one of the special configurations that are described in the topic about mixing generations of 3592 media.

Special configurations are also required for mixing different generations of 3592 drives in 349x and ACSLS libraries.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is AD5M. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional. The default is ALLOW.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes-for example, back up sets, export volumes, and database backup volumes-will not be encrypted.) If you specify ON and you enable either the library or system method of encryption, drive encryption is not allowed and backup operations fail.

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if either the library or system method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive.

When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption.

By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable either the library or system method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

DEFINE DEVCLASS (Define a 4MM device class)

Use the 4MM device class when you are using 4 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>>-LIBRary----library_name--DEVType----4MM----->
.-FORMAT----DRIVE-----
>--+-----+-----+-----+----->
'-FORMAT----+DRIVE--+' '-ESTCAPacity----size-'
      +-DDS1--+
      +-DDS1C--+
      +-DDS2--+
      +-DDS2C--+
      +-DDS3--+
      +-DDS3C--+
      +-DDS4--+
      +-DDS4C--+
      +-DDS5--+
      +-DDS5C--+
      +-DDS6--+
      '-DDS6C-'

.-PREFIX----ADSM-----
>--+-----+-----+-----+----->
'-PREFIX----+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTWait----60-----.-MOUNTRetention----60-----
>--+-----+-----+-----+----->
'-MOUNTWait----minutes-' '-MOUNTRetention----minutes-'

.-MOUNTLimit----DRIVES-----
>--+-----+-----+-----+-----><
'-MOUNTLimit----+DRIVES--+'
      +-number--+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRary (Required)

Specifies the name of the defined library object that contains the 4 mm tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=4MM (Required)

Specifies that the 4MM device type is assigned to the device class. The 4MM indicates that 4 mm tape devices are assigned to this device class.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for 4 mm devices:

Table 1. Recording formats and default estimated capacities for 4 mm tapes

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
DDS1	2.6 GB (60 meter) 4.0 GB (90 meter)	Uncompressed format, applies only to 60-meter and 90-meter tapes
DDS1C	See note 1.3 GB (60 meter) 2.0 GB (90 meter)	Compressed format, applies only to 60-meter and 90-meter tapes
DDS2	4.0 GB	Uncompressed format, applies only to 120-meter tapes
DDS2C	See note 8.0 GB	Compressed format, applies only to 120-meter tapes
DDS3	12.0 GB	Uncompressed format, applies only to 125-meter tapes
DDS3C	See note 24.0 GB	Compressed format, applies only to 125-meter tapes
DDS4	20.0 GB	Uncompressed format, applies only to 150-meter tapes
DDS4C	See note 40.0 GB	Compressed format, applies only to 150-meter tapes
DDS5	36 GB	Uncompressed format, when using DAT 72 media
DDS5C	See note 72 GB	Compressed format, when using DAT 72 media
DDS6	80 GB	Uncompressed format, when using DAT 160 media
DDS6C	See note 160 GB	Compressed format, when using DAT 160 media
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be greater than the listed value.		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

For more information about the default estimated capacity for 4 mm tapes, see Table 1

PREFIX

Specifies the high-level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DEFINE DEVCLASS (Define an 8MM device class)

Use the 8MM device class when you are using 8 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfIne DEVclAss--device_class_name----->
>--LIBRary-----library_name--DEVType-----8MM----->
  .-WORM-----No----- .-FORMAT-----DRIVE-----
>--+-----+-----+-----+-----+-----+----->
  '-WORM-----+No--+-' '-FORMAT-----+DRIVE--+-'
      '-Yes-'                +-8200--+
                          +-8200C-+
                          +-8500--+
                          +-8500C-+
                          +-8900--+
                          +-AIT---+
                          +-AITC--+
                          +-M2----+
                          +-M2C---+
                          +-SAIT--+
                          +-SAITC-+
                          +-VXA2--+
                          +-VXA2C-+
                          +-VXA3--+
                          '-VXA3C-'
>--+-----+-----+-----+-----+-----+----->
  '-ESTCAPacity-----size-'
  .-PREFIX-----ADSM-----
>--+-----+-----+-----+-----+-----+----->
  '-PREFIX-----+ADSM-----+-'
      '-tape_volume_prefix-'
  .-MOUNTRetention-----60----- .-MOUNTWait-----60-----
>--+-----+-----+-----+-----+-----+----->
  '-MOUNTRetention-----minutes-' '-MOUNTWait-----minutes-'
  .-MOUNTLimit-----DRIVES-----
>--+-----+-----+-----+-----+-----+-----><
  '-MOUNTLimit-----+DRIVES--+-'
      +-number-+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the 8 mm tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=8MM (Required)

Specifies that the 8MM device type is assigned to the device class. 8MM indicates that 8 mm tape devices are assigned to this device class.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is No. The field can contain one of the following values:

Yes

Specifies that the drives use WORM media.

No

Specifies that the drives do not use WORM media.

Note: If you select Yes, the only options available for the FORMAT parameter are:

- DRIVE
- AIT
- AITC

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for 8 mm devices:

Table 1. Recording format and default estimated capacity for 8 mm tape

Format		Description
Medium Type	Estimated Capacity	
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
8200	2.3 GB	Uncompressed (standard) format, using standard 112-meter tape cartridges
8200C	See note 3.5 GB 4.6 GB	Compressed format, using standard 112-meter tape cartridges
8500	See note	Drives (Read Write)
15m	600 MB	Eliant 820 (RW)
15m	600 MB	Exabyte 8500/8500C (RW)
15m	600 MB	Exabyte 8505 (RW)
54m	2.35 GB	Eliant 820 (RW)
54m	2.35 GB	Exabyte 8500/8500C (RW)
54m	2.35 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliant 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliant 820 (RW)

Format		Description
Medium Type	Estimated Capacity	
8500C	See note	Drives (Read Write)
15m	1.2 GB	Eliant 820 (RW)
15m	1.2 GB	Exabyte 8500/8500C (RW)
15m	1.2 GB	Exabyte 8505 (RW)
54m	4.7 GB	Eliant 820 (RW)
54m	4.7 GB	Exabyte 8500/8500C (RW)
54m	4.7 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliant 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliant 820 (RW)
8900	See note	Drive (Read Write)
15m	–	Mammoth 8900 (R)
54m	–	Mammoth 8900 (R)
112m	–	Mammoth 8900 (R)
160m XL	–	Mammoth 8900 (R)
22m	2.5 GB	Mammoth 8900 (RW)
125m	–	Mammoth 8900 (RW with upgrade)
170m	40 GB	Mammoth 8900 (RW)
AIT	See note	Drive
SDX1–25C	25 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	35 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	36 GB	AIT2 and AIT3 drives
SDX2–50C	50 GB	AIT2 and AIT3 drives
SDX3–100C	100 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	150 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	200 GB	AIT4 and AIT5 drives
SDX5-400C	400 GB	AIT5 drive
AITC	See note	Drive
SDX1–25C	50 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	91 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	72 GB	AIT2 and AIT3 drives
SDX2–50C	130 GB	AIT2 and AIT3 drives
SDX3–100C	260 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	390 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	520 GB	AIT4 and AIT5 drives
SDX5-400C	1040 GB	AIT5 drive
M2	See note	Drive (Read Write)
75m	20.0 GB	Mammoth II (RW)
150m	40.0 GB	Mammoth II (RW)
225m	60.0 GB	Mammoth II (RW)
M2C	See note	Drive (Read Write)
75m	50.0 GB	Mammoth II (RW)
150m	100.0 GB	Mammoth II (RW)
225m	150.0 GB	Mammoth II (RW)
SAIT	See note	Drive (Read Write)
	500 GB	Sony SAIT1–500(RW)
SAITC	See note	Drive (Read Write)
	1300 GB (1.3 TB)	Sony SAIT1–500(RW)

Format		Description
Medium Type	Estimated Capacity	
VXA2	See note	Drive (Read Write)
V6 (62m)	20 GB	VXA-2
V10 (124m)	40 GB	
V17 (170m)	60 GB	
VXA2C	See note	Drive (Read Write)
V6 (62m)	40 GB	VXA-2
V10 (124m)	80 GB	
V17 (170m)	120 GB	
VXA3	See note	Drive (Read Write)
X6 (62m)	40 GB	VXA-3
X10 (124m)	86 GB	
X23 (230m)	160 GB	
VXA3C	See note	Drive (Read Write)
X6 (62m)	80 GB	VXA-3
X10 (124m)	172 GB	
X23 (230m)	320 GB	
<p>Note: The actual capacities might vary depending on which cartridges and drives are used.</p> <ul style="list-style-type: none"> • For the M2C format, the normal compression ratio is 2.5:1. • For the AITC and SAITC formats, the normal compression ratio is 2.6:1. 		

ESTCAPACITY

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

For more information about the default estimated capacity for 8 mm tapes, see Table 1.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRETENTION

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

Example: Define an 8 mm device class

Define a device class that is named 8MMTAPE for an 8 mm device in a library named AUTO. The format is DRIVE, mount limit is 2, mount retention is 10, tape volume prefix is named ADSMVOL, and the estimated capacity is 6 GB.

```
define devclass 8mmtape devtype=8mm library=auto
format=drive mountlimit=2 mountretention=10
prefix=adsmvol estcapacity=6G
```

DEFINE DEVCLASS (Define a CENTERA device class)

Use the CENTERA device class when you are using EMC Centera storage devices. The CENTERA device type uses files as volumes to store data sequentially. It is similar to the FILE device class.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name--DEVType---CENTERA----->
      .-,-----
      (1)  V      |
>>-HLAddress-----ip_address+-?PEA_file----->
      .-MINCAPacity----100M-.  .-MOUNTLimit----1-----
>--+-----+-----><
      '-MINCAPacity----size-'  '-MOUNTLimit----number-'
```

Notes:

1. For each Centera device class, you must specify one or more IP addresses. However, a Pool Entry Authorization (PEA) file name and path are optional, and up to one PEA file specification can follow the IP addresses. Use the "?" character to separate the PEA file name and path from the IP addresses.

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=CENTERA (Required)

Specifies that the Centera device type is assigned to this device class. All volumes that belong to a storage pool that is defined to this device class are logical volumes that are a form of sequential access media.

HLAddress

Specifies one or more IP addresses for the Centera storage device and, optionally, the name and path of one Pool Entry Authorization (PEA) file. Specify the IP addresses with the dotted decimal format (for example, 9.10.111.222). A Centera device might have multiple IP addresses. If multiple IP addresses are specified, then the store or retrieve operation attempts a connection by using each IP address that is specified until a valid address is found.

AIX The PEA file name and path name are case-sensitive.

If you append the name and path of a PEA file, ensure that the file is stored in a directory on the system that runs the server. Separate the PEA file name and path from the IP address with the "?" character, for example: **Windows**

```
HLADDRESS=9.10.111.222,9.10.111.223?c:\controlFiles\TSM.PEA
```

AIX

```
HLADDRESS=9.10.111.222,9.10.111.223?/user/ControlFiles/TSM.PEA
```

Specify only one PEA file name and path for each device class definition. If you specify two different Centera device classes that point to the same Centera storage device and if the device class definitions contain different PEA file names and paths, the server uses the PEA file that is specified in the device class HLADDRESS parameter that was first used to open the Centera storage device.

Tips:

1. The server does not include a PEA file during installation. If you do not create a PEA file, the server uses the Centera default profile, which can allow applications to read, write, delete, purge, and query data on a Centera storage device. To provide tighter control, create a PEA file with the command-line interface that is provided by EMC Centera. For details about Centera authentication and authorization, refer to the EMC Centera *Programmer's Guide*.
2. You can also specify the PEA file name and path in an environment variable with the syntax `CENTERA_PEA_LOCATION=filePath_fileName`. The PEA file name and path that is specified with this environment variable apply to all Centera clusters. If you use this variable, you do not have to specify the PEA file name and path with the HLADDRESS parameter.

MINCAPacity

Specifies the minimum size for Centera volumes that are assigned to a storage pool in this device class. This value represents the minimum amount of data that is stored on a Centera volume before the server marks it full. Centera volumes continue to accept data until the minimum amount of data is stored. This parameter is optional.

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The default value is 100 MB (MINCAPACITY=100M). The minimum value that is allowed is 1 MB (MINCAPACITY=1M). The maximum value that is allowed is 128 GB (MINCAPACITY=128G).

MOUNTLimit

Specifies the maximum number of files that can be simultaneously open for input and output. The default value is 1. This parameter is optional. You can specify any number from 0 or greater; however, the sum of all mount limit values for all device classes that are assigned to the same Centera device must not exceed the maximum number of sessions that are allowed by Centera.

DEFINE DEVCLASS (Define a DLT device class)

Use the DLT device class when you are using DLT tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRARY----library_name--DEVType----DLT----->
.-WORM----No----- .-FORMAT----DRIVE-----
>--+-----+-----+-----+----->
'-WORM----+-No-+-' '-FORMAT----+-DRIVE----+'
      '-Yes-'                +-DLT1-----+
                          +-DLT1C----+
                          +-DLT10----+
                          +-DLT10C---+
                          +-DLT15----+
                          +-DLT15C---+
                          +-DLT20----+
                          +-DLT20C---+
                          +-DLT35----+
                          +-DLT35C---+
                          +-DLT40----+
                          +-DLT40C---+
                          +-DLT2-----+
                          +-DLT2C----+
                          +-DLT4-----+
                          +-DLT4C----+
                          +-SDLT-----+
                          +-SDLTC----+
                          +-SDLT320--+
                          +-SDLT320C--+
                          +-SDLT600--+
                          +-SDLT600C--+
                          +-DLTS4-----+
                          '-DLTS4C---'
```

```
>--+-----+-----+-----+----->
'-ESTCAPacity----size-'
.-PREFIX----ADSM-----
>--+-----+-----+-----+----->
'-PREFIX----+-ADSM-----+'
      '-tape_volume_prefix-'
```

```
.-MOUNTRetention----60----- .-MOUNTWait----60-----
>--+-----+-----+-----+----->
'-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'
```

```
.-MOUNTLimit----DRIVES-----
>--+-----+-----+-----+-----><
'-MOUNTLimit----+-DRIVES-+-'
      +-number-+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the DLT tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=DLT (Required)

Specifies that the DLT device type is assigned to the device class. DLT indicates that DLT tape devices are assigned to this device class.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is No. The field can contain one of the following values:

Yes

Specifies that the drives use WORM media.

No

Specifies that the drives do not use WORM media.

Note: Support for DLT WORM media is available only for SDLT-600, Quantum DLT-V4, and Quantum DLT-S4 drives in manual, SCSI, and ACSLS libraries.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for DLT devices:

Table 1. Recording format and default estimated capacity for DLT

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
DLT1	40.0 GB	Uncompressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT1C	See note 1. 80.0 GB	Compressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT10	10.0 GB	Uncompressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT10C	See note 1. 20.0 GB	Compressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT15	15.0 GB	Uncompressed format, using only CompacTape IIIxt cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT15C	See note 1. 30.0 GB	Compressed format, using only CompacTape IIIxt cartridges Valid with DLT4000, DLT7000, and DLT8000 drives

Format	Estimated Capacity	Description
DLT20	20.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT20C	See note 1. 40.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT35	35.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT35C	See note 1. 70.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT40	40.0 GB	Uncompressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT40C	See note 1. 80.0 GB	Compressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT2	80.0 GB	Uncompressed format, using Quantum DLT tape VS1 media
DLT2C	See note 1. 160.0 GB	Compressed format, using Quantum DLT tape VS1 media
DLT4	160.0 GB	Uncompressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
DLT4C	See note 1. 320.0 GB	Compressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
SDLT See note 2.	100.0 GB	Uncompressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLTC See note 2.	See note 1. 200.0 GB	Compressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLT320 See note 2.	160.0 GB	Uncompressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT320C See note 2.	See note 1. 320.0 GB	Compressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT600	300.0 GB	Uncompressed format, using SuperDLTtape-II media Valid with a Super DLT drive
SDLT600C	See note 1. 600.0 GB	Compressed format, using SuperDLTtape-II media Valid with a Super DLT drive
DLTS4	800 GB	Uncompressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive
DLTS4C	See note 1. 1.6 TB	Compressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive

Format	Estimated Capacity	Description
<p>Note:</p> <ol style="list-style-type: none"> 1. Depending on the effectiveness of compression, the actual capacity might be greater than the listed value. 2. IBM Spectrum Protect™ does not support a library that contains both Backward Read Compatible (BRC) SDLT and Non-Backward Read Compatible (NBRC) SDLT drives. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

For more information about estimated capacities, see Table 1.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DEFINE DEVCLASS (Define an ECARTRIDGE device class)

Use the ECARTRIDGE device class when you are using StorageTek drives such as the StorageTek T9840 or T10000.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>>-LIBRARY-----library_name--DEVType-----ECARTridge----->
                                     (1)
.-LBProtect-----No----- .-WORM-----No-----
>>+-----+-----+-----+----->
'-LBProtect-----+READWrite+-' '-WORM-----+No--+-'
      +WRITEOnly+          '-Yes-'
      '-No-----'
.-FORMAT-----DRIVE-----
>>+-----+-----+-----+----->
'-FORMAT-----+DRIVE-----+' '-ESTCAPacity-----size-'
      +T9840C----+
      +T9840C-C--+
      +T9840D----+
      +T9840D-C--+
      +T10000A---+
      +T10000A-C+
      +T10000B---+
      +T10000B-C+
      +T10000C---+
      +T10000C-C+
      +T10000D---+
      '-T10000D-C-'
.-PREFIX-----ADSM-----
>>+-----+-----+-----+----->
'-PREFIX-----+ADSM-----+'
      '-tape_volume_prefix-'
.-MOUNTRetention-----60----- .-MOUNTWait-----60-----
>>+-----+-----+-----+----->
'-MOUNTRetention-----minutes-' '-MOUNTWait-----minutes-'
```

```

.-MOUNTLimit-----DRIVES-----
>-----+-----+-----+-----+-----+----->
'-MOUNTLimit-----+DRIVES-+-'
      +-number-+
      '-0-----'

(1) (2)
.-DRIVEEncryption-----ALLOW-----
>-----+-----+-----+-----+-----+----->>
'-DRIVEEncryption-----+ON-----+--'
      +-ALLOW-----+
      +-EXternal-+
      '-OFF-----'

```

Notes:

1. You cannot specify both WORM=Yes and DRIVEENCRYPTION=ON.
2. You can use drive encryption only for Oracle StorageTek T10000B drives with a format value of DRIVE, T10000B, or T10000B-C, for Oracle StorageTek T10000C drives with a format value of DRIVE, T10000C or T10000C-C, and for Oracle StorageTek T10000D drives with a format value of DRIVE, T10000D and T10000D-C.

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the ECARTRIDGE tape drives that can be used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=ECARtridge (Required)

Specifies that the ECARTRIDGE device type is assigned to the device class. ECARTRIDGE indicates that a specific type of cartridge tape device (StorageTek) is assigned to this device class.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When LBPROTECT is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is NO.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect™ and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the BACKUP DB command.

When the LBPROTECT parameter is set to READWRITE, you do not have to specify the CRCDATA parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the BACKUP DB command.

No

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on Oracle StorageTek T10000C and Oracle StorageTek T10000D drives.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is No. The field can contain one of the following values:

Yes

Specifies that the drives use WORM media.

No

Specifies that the drives do not use WORM media.

Restriction: If you select Yes, the only options that are available for the FORMAT parameter are:

- DRIVE
- T9840C
- T9840C-C
- T9840D
- T9840D-C
- T10000A
- T10000A-C
- T10000B
- T10000B-C
- T10000C
- T10000C-C
- T10000D
- T10000D-C

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

Important: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format that is established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use. The following table lists the recording formats and estimated capacities for ECARTRIDGE devices:

Table 1. Recording formats and default estimated capacities for ECARTRIDGE tapes

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge

Format	Estimated capacity	Description
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
<p>Notes:</p> <ul style="list-style-type: none"> • Some formats use a tape drive hardware compression feature. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value. • T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats. T10000D drives can read, but cannot write, the T10000A, T10000B, and T10000C formats. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is AD5M. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB . CD2 . E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is AD5M.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional. The default is ALLOW.

Restrictions:

1. You can use drive encryption only for the following drives:
 - o Oracle StorageTek T10000B drives that have a format value of DRIVE, T10000B, or T10000B-C
 - o Oracle StorageTek T10000C drives that have a format value of DRIVE, T10000C, or T10000C-C
 - o Oracle StorageTek T10000D drives that have a format value of DRIVE, T10000D, or T10000D-C
2. You cannot specify IBM Spectrum Protect as the key manager for drive encryption of write once, read many (WORM) media. You cannot specify both WORM=Yes and DRIVEENCRYPTION=ON.
3. If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, back up sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not allowed and backup operations fail.

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if another method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption. By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable another method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

DEFINE DEVCLASS (Define a FILE device class)

Use the FILE device class when you are using files on magnetic disk storage as volumes that store data sequentially (as on tape).

AIX | **Linux** The FILE device class does not support EXTERNAL libraries.

Windows The FILE device class does not support EXTERNAL or Remote Storage Manager libraries.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see DEFINE DEVCLASS (Define a FILE device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfINE DEVclAss--device_class_name--DEVType==--FILE----->
      .-MOUNTLimit----20----- .-MAXCAPacity----10G--.
>+-----+-----+-----+-----+-----+-----+----->
      '-MOUNTLimit----number-' '-MAXCAPacity----size-'
      .-DIRectory----current_directory_name-.
>+-----+-----+-----+-----+-----+-----+----->
      |           .-,'-----'          |
      |           v           |          |
      '-DIRectory----directory_name-+-----'
      .-SHAREd----No----- .
>+-----+-----+-----+-----+-----+-----+-----><
      '-SHAREd----+No--+-'
              '-Yes-'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=FILE (Required)

Specifies that the FILE device type is assigned to the device class. FILE indicates that a file is assigned to this device class. When the server must access a volume that belongs to this device class, it opens a file and reads or writes file data.

A file is a form of sequential-access media.

MOUNTLimit

Specifies the maximum number of files that can be simultaneously open for input and output. This parameter is optional. The default value is 20. You can specify a number from 0 to 4096.

Windows If the device class is shared with a storage agent (by specifying the SHARED=YES parameter), drives are defined or deleted to match the mount limit value.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

MAXCAPacity

Specifies the maximum size of any data storage files that are defined to a storage pool in this device class.

The value of the MAXCAPACITY parameter is also used as the unit of allocation when storage pool space triggers create volumes. The default value is 10 GB (MAXCAPACITY=10G). The value that is specified must be less than or equal to the

maximum supported size of a file on the target file system.

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The minimum size is 1 MB (MAXCAPACITY=1M). If you are defining a FILE device class for database-backup volumes, specify a value for MAXCAPACITY that is appropriate for the size of the database and that minimizes the number of database volumes.

AIX | **Linux** Do not define a MAXCAPACITY value greater than 640M when this file is for REMOVABLEFILE CD support. A value less than a CD's usable space (650 MB) enables a one-to-one match between files from the FILE device class and copies that are on CD.

DIRECTORY

Specifies the directory location or locations of the files that are used in this device class. Enclose the entire list of directories within quotation marks, and use commas to separate individual directory names. Special characters (for example, blank spaces) are allowed within directory names. For example, the directory list "abc def,xyz" contains two directories: abc def and xyz.

This parameter is optional.

AIX | **Linux** The default is the current working directory of the server at the time the command is issued.

Windows The default is the current working directory of the server at the time the command is issued. Windows registry information is used to determine the default directory.

By specifying a directory name or names, you identify the location where the server places the files that represent storage volumes for this device class.

For NetApp SnapLock support (storage pools with RECLAMATIONTYPE=SNAPLOCK, which are going to use this device class), the directory, or directories that are specified with DIRECTORY parameter must point to the directory or directories on the NetApp SnapLock volumes.

AIX | **Linux** While the command is processed, the server expands the specified directory name or names into their fully qualified forms, starting from the root directory.

If the server must allocate a scratch volume, it creates a new file in one of these directories. (The server can choose any of the directories in which to create new scratch volumes.) For scratch volumes used to store client data, the file that is created by the server has a file name extension of .bfs. For scratch volumes used to store export data, a file name extension of .exp is used.

AIX | **Linux** For example, if you define a device class with a directory of tsmstor and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named tsmstor\00566497.exp.

Windows For example, if you define a device class with a directory of c:\server and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named c:\server\00566497.exp.

Important: You must ensure that storage agents can access newly created FILE volumes. Failure of the storage agent to access a FILE volume can cause operations to be retried on a LAN-only path or to fail. For more information, see the description of the DIRECTORY parameter in DEFINE PATH (Define a path).

Tip: If you specify multiple directories for a device class, ensure that the directories are associated with separate file systems. Space trigger functions and storage pool space calculations take into account the space that remains in each directory. If you specify multiple directories for a device class and the directories are in the same file system, the server calculates space by adding values that represent the space that remains in each directory. These space calculations are inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the wrong storage pool and run out of space prematurely. For space triggers, an inaccurate calculation might result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled. If a trigger is disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by issuing the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

SHARED

Specifies that this FILE device class is shared between the server and one or more storage agents. To prepare for sharing, a library is automatically defined along with a number of drives corresponding to the MOUNTLIMIT parameter value. The drive names are the name of the library plus a number from 1 to the mount limit number. For example, if the library name is FILE and the mount limit is set to 4, the drives are named FILE11, FILE12, FILE13, FILE14.

For information about prerequisites when storage is shared by the server and storage agent, see IBM® Support Portal for IBM Spectrum Protect™.

Example: Define a FILE device class with multiple directories

Define a device class that specifies multiple directories.

AIX

```
define devclass multidir devtype=file
  directory=/usr/xyz,/usr/abc,/usr/uvw
```

Linux

```
define devclass multidir devtype=file
  directory=/opt/xyz,/opt/abc,/opt/uvw
```

Windows

```
define devclass multidir devtype=file
  directory=e:\xyz,f:\abc,g:\uvw
```

Example: Define a FILE device class with a 50 MB capacity

Define a device class named PLAINFILES with a FILE device type and a maximum capacity of 50 MB.

```
define devclass plainfiles devtype=file
maxcapacity=50m
```

AIX

Windows

DEFINE DEVCLASS (Define a GENERICTAPE device class)

Use the GENERICTAPE device class for tape drives that are supported by operating system device drivers.

When you use this device type, the server does not recognize either the type of device or the cartridge recording format. Because the server does not recognize the type of device, if an I/O error occurs, error information is less detailed compared to error information for a specific device type (for example, 8MM). When you define devices to the server, do not mix various types of devices within the same device type.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRARY----library_name--DEVType----GENERICtape----->
                                     .-MOUNTRetention----60-----.
>--+-----+-----+-----+-----+----->
  '-ESTCAPacity----size-'  '-MOUNTRetention----minutes-'
                                     .-MOUNTWait----60-----.  .-MOUNTLimit----DRIVES-----.
>--+-----+-----+-----+-----+-----><
  '-MOUNTWait----minutes-'  '-MOUNTLimit----+DRIVES+-'
                                     +-number+
                                     '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the tape drives that can be used by this device class.

For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=GENERICtape (Required)

Specifies that the GENERICTAPE device type is assigned to the device class. GENERICTAPE indicates that the volumes for this device class are used in tape drives that are supported by the operating system's tape device driver.

The server recognizes that the media can be removed and that more media can be inserted, subject to limits set with the MOUNTLIMIT parameter for the device class and the MAXSCRATCH parameter for the storage pool.

Volumes in a device class with device type GENERICTAPE are sequential access volumes.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

Specify a capacity appropriate to the particular tape drive that is being used.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DEFINE DEVCLASS (Define an LTO device class)

Use the LTO device class when you are using LTO tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>>-LIBRARY----library_name--DEVType----LTO----->

      (1)
.-LBProtect----No----- .-WORM----No-----
>--+-----+-----+-----+----->
'-LBProtect----+READWrite+-' '-WORM----+No--+'
      +-WRITEOnly+          '-Yes-'
      '-No-----'

.-FORMAT----DRIVE-----
>--+-----+-----+-----+----->
|          (2)          | '-ESTCAPacity----size-'
'-FORMAT----+DRIVE----+'
      +-ULTRIUM2---+
      +-ULTRIUM2C--+
      +-ULTRIUM3---+
      +-ULTRIUM3C--+
      +-ULTRIUM4---+
      +-ULTRIUM4C--+
      +-ULTRIUM5---+
      +-ULTRIUM5C--+
      +-ULTRIUM6---+
      +-ULTRIUM6C--+
      +-ULTRIUM7---+
      +-ULTRIUM7C--+
      +-ULTRIUM8---+
      '-ULTRIUM8C-'

.-PREFIX----ADSM-----
>--+-----+-----+-----+----->
'-PREFIX----+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTRetention----60----- .-MOUNTWait----60-----
>--+-----+-----+-----+----->
'-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'

.-MOUNTLimit----DRIVES-----
>--+-----+-----+-----+----->
'-MOUNTLimit----+DRIVES--+'
      +-number-+
      '-0-----'

      (1) (3)
.-DRIVEEncryption----ALLOW-----
>--+-----+-----+-----+-----><
'-DRIVEEncryption----+ON-----+'
      +-ALLOW-----+
      +-EXTERNAL--+
      '-OFF-----'
```

Notes:

1. You cannot specify both WORM=Yes and DRIVEENCRYPTION=ON.
2. IBM Spectrum Protect™ server supports LTO-2 tape drives; however, IBM® Tape Device drivers do not. In the event of an issue with the LTO-2 drive, the preferred corrective action is to upgrade your tape drive hardware to a higher generation

drive, then install the latest version of the device driver.

3. Drive encryption is supported only for LTO-4 and higher generation LTO drives and media.

Parameters

`device_class_name` (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

`LIBRARY` (Required)

Specifies the name of the defined library object that contains the LTO tape drives used by this device class. For information about defining a library object, see the `DEFINE LIBRARY` command.

`DEVType=LTO` (Required)

Specifies that the linear tape open (LTO) device type is assigned to the device class.

`LBProtect`

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When `LBPROTECT` is set to `READWRITE` or to `WRITEONLY`, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is `NO`.

The following values are possible:

`READWrite`

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect and the tape drive to calculate and compare CRC values. The `READWRITE` value does not affect backup sets and data that is generated by the `BACKUP DB` command.

When the `LBPROTECT` parameter is set to `READWRITE`, you do not have to specify the `CRCDATA` parameter in a storage pool definition because logical block protection provides better protection against data corruption.

`WRITEOnly`

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The `WRITEONLY` value does not affect backup sets and data that are generated by the `BACKUP DB` command.

`No`

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction:

Restrictions apply to logical block protection (LBP):

- At the LTO-5 level, LBP is supported only on IBM LTO-5.
- Starting with LTO-6, LBP is supported by all LTO drive vendors.

`WORM`

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is `No`. The field can contain one of the following values:

`Yes`

Specifies that the drives use WORM media.

`No`

Specifies that the drives do not use WORM media.

Note:

1. To use WORM media in a library, all the drives in the library must be WORM capable.
2. You cannot specify IBM Spectrum Protect as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both `WORM=Yes` and `DRIVEENCRYPTION=ON` is not supported.)

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

If you are considering mixing different generations of LTO media and drives, be aware of the following restrictions.

Table 1. Read - write capabilities for different generations of LTO drives

Drives	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media	Generation 7 media	Generation M8 media	Generation 8 media
Generation 3 ¹	Read and write	n/a	n/a	n/a	n/a	n/a	n/a
Generation 4 ¹	Read and write	Read and write	n/a	n/a	n/a	n/a	n/a
Generation 5 ¹	Read only	Read and write	Read and write	n/a	n/a	n/a	n/a
Generation 6 ¹	n/a	Read only	Read and write	Read and write	n/a	n/a	n/a
Generation 7 ¹			Read only	Read and write	Read and write	n/a	n/a
Generation 8 ²	n/a	n/a	n/a	n/a	Read and write	Read and write	Read and write

¹ If a storage pool volume can only be read by a tape drive, ensure that the attributes of the storage pool volume are set to read only.

² LTO-8 drives have two media types: LTO-M8 media and LTO-8 media. Both media types are used only in LTO-8 tape drives.

The following table lists the recording formats and estimated capacities for LTO devices:

Table 2. Recording format and default estimated capacity for LTO

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
ULTRIUM2	200 GB	Uncompressed (standard) format, using Ultrium 2 cartridges
ULTRIUM2C	See note 400 GB	Compressed format, using Ultrium 2 cartridges
ULTRIUM3	400 GB	Uncompressed (standard) format, using Ultrium 3 cartridges
ULTRIUM3C	See note 800 GB	Compressed format, using Ultrium 3 cartridges
ULTRIUM4	800 GB	Uncompressed (standard) format, using Ultrium 4 cartridges

Format	Estimated capacity	Description
ULTRIUM4C	See note 1.6 TB	Compressed format, using Ultrium 4 cartridges
ULTRIUM5	1.5 TB	Uncompressed (standard) format, using Ultrium 5 cartridges
ULTRIUM5C	Varied, as described in note	Compressed format, using Ultrium 5 cartridges
ULTRIUM6	2.5 TB	Uncompressed (standard) format, using Ultrium 6 cartridges
ULTRIUM6C	Varied, as described in note	Compressed format, using Ultrium 6 cartridges
ULTRIUM7	6 TB	Uncompressed (standard) format, using Ultrium 7 cartridges
ULTRIUM7C	Varied, as described in note	Compressed format, using Ultrium 7 cartridges
ULTRIUM8	12 TB for LTO-8 media 9 TB for LTO-M8 media	Uncompressed (standard) format, using Ultrium M8 or Ultrium 8 cartridges
ULTRIUM8C	Varied, as described in note	Compressed format, using Ultrium M8 or Ultrium 8 cartridges
Note: If this format uses the tape-drive hardware-compression feature, depending on the effectiveness of compression, the actual capacity is varied.		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

For more information about estimated capacities, see Table 2.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional. The default is ALLOW. Drive encryption is supported only for LTO-4 and higher generation drives and media.

Restriction: If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, back up sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not allowed and backup operations fail.

Note: You cannot specify IBM Spectrum Protect as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both WORM=Yes and DRIVEENCRYPTION=ON is not supported.)

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if another method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption. By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable another method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

Example: Define an LTO device class

Define a device class that is named LTOTAPE for an LTO drive in a library named LTOLIB. The format is ULTRIUM, mount limit is 12, mount retention is 5, tape volume prefix is named SMVOL, and the estimated capacity is 100 GB.

```
define devclass ltotape devtype=lto library=ltolib
format=ultrium mountlimit=12 mountretention=5
prefix=smvol estcapacity=100G
```

DEFINE DEVCLASS (Define a NAS device class)

Use the NAS device class when you are using NDMP (Network Data Management Protocol) operations to back up network-attached storage (NAS) file servers. The device class is for drives that are supported by the NAS file server for backups.

AIX | **Linux** The NAS device class does not support EXTERNAL libraries.

Windows The NAS device class does not support EXTERNAL or Remote Storage Manager libraries.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name--DEVType--==--NAS----->
>--LIBRARY-----library_name--MOUNTRetention-----0----->
    .-MOUNTWait-----60----- .-MOUNTLimit-----DRIVES-----
>--+-----+-----+-----+-----+----->
    '-MOUNTWait-----minutes-' '-MOUNTLimit-----+DRIVES--+'
                                     +-number+
                                     '-0-----'

>--ESTCAPacity-----size----->
    .-PREFIX-----ADSM-----
>--+-----+-----+-----+-----+----->>
    '-PREFIX-----+ADSM-----+'
        '-tape_volume_prefix-'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=NAS (Required)

Specifies that the network-attached storage (NAS) device type is assigned to the device class. The NAS device type is for drives that are attached to and used by a NAS file server for backup of NAS file systems.

LIBRARY (Required)

Specifies the name of the defined library object that contains the SCSI tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

MOUNTRetention=0 (Required)

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. Zero (0) is the only supported value for device classes with DEVType=NAS.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

ESTCAPacity (Required)

Specifies the estimated capacity for the volumes that are assigned to this device class.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

Example: Define a NAS device class

Define a device class that is named NASTAPE for a NAS drive in a library named NASLIB. The mount limit is DRIVES, mount retention is 0, tape volume prefix is named SMVOL, and the estimated capacity is 200 GB.

```
define devclass nastape devtype=nas library=naslib
mountretention=0 mountlimit=drives
prefix=smvol estcapacity=200G
```

DEFINE DEVCLASS (Define a REMOVABLEFILE device class)

Use the REMOVABLEFILE device class for removable media devices that are attached as local, removable file systems.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRary----library_name--DEVType----REMOVABLEfile----->
  .-MAXCAPacity----space_remaining-.
>--+-----+-----+----->
  '-MAXCAPacity----size-----'
  .-MOUNTRetention----60----- .-MOUNTWait----60-----
>--+-----+-----+----->
  '-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'
  .-MOUNTLimit----DRIVES-----
>--+-----+-----+-----><
  '-MOUNTLimit----+DRIVES++'
      +-number-+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the removable media drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=REMOVABLEfile (Required)

Specifies that the REMOVABLEFILE device type is assigned to the device class. REMOVABLEFILE indicates that the volumes for this device class are files on local, removable media.

Volumes in a device class with device type REMOVABLEFILE are sequential access volumes.

Use the device manufacturer's utilities to format (if necessary) and label the media. The label on the media must meet the following restrictions:

- The label can have no more than 11 characters.
- The volume label and the name of the file on the volume must match exactly.
- **AIX** | **Windows** The MAXCAPACITY parameter value must be specified at less than the capacity of the media.

MAXCAPacity

Specifies the maximum size of any volumes that are defined to a storage pool categorized by this device class. This parameter is optional.

The MAXCAPACITY parameter must be set at less value than the capacity of the media. For CD media, the maximum capacity can be no greater than 650 MB.

AIX | **Windows** Because the server opens only one file per physical removable medium, specify a capacity that enables one file to make full use of your media capacity.

space_remaining

The default maximum capacity is the space that remains on the media after it is first used.

size

You must specify this value as an integer followed by K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes).

For example, MAXCAPACITY=5M specifies that the maximum capacity for a volume in this device class is 5 MB. The smallest value that is allowed is 1 MB (that is, MAXCAPACITY=1M).

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DEFINE DEVCLASS (Define a SERVER device class)

Use the SERVER device class to use storage volumes or files that are archived in another IBM Spectrum Protect™ server.

If data retention protection is activated with the SET ARCHIVERETENTIONPROTECTION command, you cannot define a server device class.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name--DEVType---SERVER----->
                                     .-MAXCAPacity---500M-.
>--SERVERName---server_name--+-----+----->
                                     '-MAXCAPacity---size-'

    .-MOUNTLimit---1----- .-MOUNTRetention---60-----
>--+-----+-----+-----+----->
    '-MOUNTLimit---number-' '-MOUNTRetention---minutes-'

    .-PREFIX---ADSM-----
>--+-----+-----+-----+----->
    '-PREFIX---+ADSM-----+-'
                                     '-volume_prefix-'
```

```

.-RETRYPeriod----10-----
>-----+----->
'-RETRYPeriod----retry_value_(minutes)-'

.-RETRYInterval----30-----
>-----+----->>
'-RETRYInterval----retry_value_(seconds)-'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=SERVER (Required)

Specifies a remote connection that supports virtual volumes.

SERVERName (Required)

Specifies the name of the server. The SERVERNAME parameter must match a defined server.

MAXCAPacity

Specifies the maximum size for objects that are created on the target server; the default for this value is 500M. This parameter is optional.

500M

Specifies that the maximum capacity is 500M (500 MB).

size

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The minimum value that is allowed is 1 MB (MAXCAPACITY=1M).

MOUNTLimit

Specifies the maximum number of simultaneous sessions between the source server and the target server. Any attempts to access more sessions than indicated by the mount limit cause the requester to wait. This parameter is optional. The default value is 1. You can specify a number 1 - 4096.

The following are possible values:

1

Specifies that only one session between the source server and the target server is allowed.

number

Specifies the number of simultaneous sessions between the source server and the target server.

MOUNTRetention

Specifies the number of minutes to retain an idle connection with the target server before the connection closes. This parameter is optional. The default value is 60. You can specify a number 0 - 9999.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The default is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high-level archive file name that uses the default prefix is ADSM.volume1.

RETRYPeriod

Specifies the retry period in minutes. The retry period is the interval during which the server attempts to contact a target server if there is a suspected communications failure. This parameter is optional. You can specify a number 0 - 9999. The default value is 10 minutes.

RETRYInterval

Specifies the retry interval in seconds. The retry interval is how often retries are done within a specific time period. This parameter is optional. You can specify a number 1 - 9999. The default value is 30 seconds.

DEFINE DEVCLASS (Define a VOLSAFE device class)

Use the VOLSAFE device type to work with StorageTek VolSafe brand media and drives. This technology uses media that cannot be overwritten. Therefore, do not use these media for short-term backups of client files, the server database, or export tapes.

Restrictions:

1. NAS-attached libraries are not supported.
2. VolSafe media and read/write media must be in separate storage pools.
3. Check in cartridges with CHECKLABEL=YES on the CHECKIN LIBVOLUME command.
4. Label cartridges with OVERWRITE=NO on the LABEL LIBVOLUME command. If VolSafe cartridges are labeled more than one time, no additional data can be written to them.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfIne DEVclass--device_class_name----->
>>-LIBRARY----library_name--DEVType----VOLSAFE----->
>>-WORM----Yes-----+----->
      .-FORMAT----DRIVE-----
      '-FORMAT----+DRIVE-----+'
          +-9840-----+
          +-9840-C----+
          +-T9840C----+
          +-T9840C-C--+
          +-T9840D----+
          +-T9840D-C--+
          +-T10000A---+
          +-T10000A-C+
          +-T10000B---+
          +-T10000B-C+
          +-T10000C---+
          +-T10000C-C+
          +-T10000D---+
          '-T10000D-C-'
      .-MOUNTRetention----60-----
>+-----+----->
      '-ESTCAPacity----size-' '-MOUNTRetention----minutes-'
      .-PREFIX----ADSM-----
>+-----+----->
      '-PREFIX----+ADSM-----+'
          '-volume_prefix-'
      .-MOUNTWait----60----- .-MOUNTLimit----DRIVES-----
>+-----+-----+-----><
      '-MOUNTWait----minutes-' '-MOUNTLimit----+DRIVES++-'
          +-number+
          '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the VolSafe drives that can be used by this device class. If any drives in a library are VolSafe-enabled, all drives in the library must be VolSafe-enabled. Consult your hardware documentation to enable VolSafe on the 9840 and T10000 drives.

For information about defining a library object, see DEFINE LIBRARY (Define a library).

DEVType=VOLSAFE (Required)

Specifies that the VOLSAFE device type is assigned to the device class. The label on this type of cartridge can be overwritten one time, which IBM Spectrum Protect™ does when it writes the first block of data. Therefore, it is important to limit the use of the LABEL LIBVOLUME command to one time per volume by using the OVERWRITE=NO parameter.

WORM

Specifies whether the drives use WORM (write once, read many) media. The parameter is required. The value must be Yes.

Yes

Specifies that the drives use WORM media.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

Important: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format that is established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for VolSafe devices:

Table 1. Recording formats and default estimated capacities for Volsafe media

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
9840	20 GB	Uncompressed (standard) format, using a 20 GB cartridge with 270 meters (885 feet) of tape
9840-C	See note 80 GB	LZ-1 Enhanced (4:1) compressed format, using an 80 GB cartridge with 270 meters (885 feet) of tape
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge

Format	Estimated Capacity	Description
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

For more information about the default estimated capacity for cartridge tapes, see Table 1.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. The default value is 60 minutes. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The default is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high-level archive file name that uses the default prefix is ADSM.volume1.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.



DEFINE DEVCLASS - z/OS media server (Define device class for z/OS media server)

Use the DEFINE DEVCLASS command to define a device class for a type of storage device. The server requires that a device class be defined to allow the use of a device. A limited set of device class types is available for devices that are accessed through a z/OS® media server.

- DEFINE DEVCLASS (Define a 3590 device class for z/OS media server)
- DEFINE DEVCLASS (Define a 3592 device class for z/OS media server)
- DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server)
- DEFINE DEVCLASS (Define a FILE device class for z/OS media server)

Table 1. Commands related to DEFINE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up IBM Spectrum Protect device information to a file.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
UPDATE DEVCLASS (z/OS media server)	Changes the attributes of a device class for storage managed by a z/OS media server.



DEFINE DEVCLASS (Define a 3590 device class for z/OS media server)

To use a z/OS® media server to access 3590 devices, you must define a 3590 device class. In the device class definition, specify a library that was defined with the LIBTYPE=ZOSMEDIA parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>--LIBRARY----zos_media_library--DEVType----3590----->
                                     .-ESTCAPacity---9G-----.
>--+-----+-----+-----+-----+-----+-----+----->
  '-FORMAT---+DRIVE---+' '-ESTCAPacity-----size---'
      +-3590B---+
      +-3590C---+
      +-3590E-B-+
```

```

++3590E-C++
++3590H-B++
'-3590H-C-'

.-PREFIX---ADSM-----
>-----+-----+-----+-----+-----+-----+----->
'-PREFIX---+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTRetention---60-----.  .-MOUNTWait---60-----.
>-----+-----+-----+-----+-----+-----+----->
'-MOUNTRetention---minutes-'  '-MOUNTWait---minutes-'

.-MOUNTLimit---2-----.  .-COMpression---Yes-----.
>-----+-----+-----+-----+-----+-----+----->
'-MOUNTLimit---+DRIVES-+-'  '-COMpression---+Yes-+-'
      +-number-+                '-No--'
      '-0-----'

>-----+-----+-----+-----+-----+-----+----->
+-EXPIration---yyyddd+
'-RETention---days-----'

.-PROtection---No-----.  .-UNIT---3590-----.
>-----+-----+-----+-----+-----+-----+----->>
'-PROtection---+No-----+'  '-UNIT---unit_name-'
      +-Yes-----+
      '-Automatic-'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the DEFINE LIBRARY command.

DEVtype=3590 (Required)

Specifies the 3590 device type is assigned to the device class. 3590 indicates that 3590 cartridge tape devices are assigned to the device class.

Restriction: The z/OS media server supports 256 KB data blocks when writing to 3590 tape drives. Verify that your hardware supports this capability.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. See the following table for the recording formats.

Table 1. Recording formats for 3590

Format	Description
3590B	Uncompressed (basic) format
3590C	Compressed format
3590E-B	Uncompressed (basic) format, similar to the 3590B format
3590E-C	Compressed format, similar to the 3590C format
3590H-B	Uncompressed (basic) format, similar to the 3590B format
3590H-C	Compressed format, similar to the 3590C format
Note: If the format uses the tape drive hardware compression feature the actual capacity can increase, depending on the effectiveness of compression.	

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional. The default estimated capacity for 3590 tapes is 9 GB.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: **K** (KB), **M** (MB), **G** (GB), or **T** (TB). For example, specify that the estimated capacity is 9 GB with the parameter `ESTCAPACITY=9G`. The smallest value that is accepted is 100 KB (`ESTCAPACITY=100K`).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is `ADSM`. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

`AB.CD2.E`

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is `ADSM.BFS`.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (`LIBTYPE=EXTERNAL`), do not specify the `MOUNTWAIT` parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the `MOUNTLIMIT` parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is YES.

You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as *2014007* (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3590 tape. This parameter is optional. The default unit name is 3590. The unit name can be up to 8 characters.

AIX Linux

DEFINE DEVCLASS (Define a 3592 device class for z/OS media server)

To use a z/OS® media server to access 3592 devices, you must define a 3592 device class. In the device class definition, specify a library that was defined with the LIBTYPE=ZOSMEDIA parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>>-LIBRARY----zos_media_library--DEVType----3592----->
.-FORMAT----Drive-----.-WORM----No-----.
>--+-----+-----+----->
'-FORMAT----+DRIVE---+' '-WORM----+Yes-+'
      +-3592----+          '-No--'
      +-3592C---+
      +-3592-2---+
      +-3592-2C--+
      +-3592-3---+
      +-3592-3C--+
      +-3592-4---+
      '-3592-4C-'

.-ESTCAPacity----300G-.
>--+-----+-----+----->
'-ESTCAPacity----size-'

.-PREFIX----ADSM-----.
>--+-----+-----+----->
'-PREFIX----+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTRetention----60-----.-MOUNTWait----60-----.
>--+-----+-----+----->
'-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'

.-MOUNTLimit----2-----.-COMPRESSION----Yes-----.
>--+-----+-----+----->
'-MOUNTLimit----+DRIVES-+' '-COMPRESSION----+Yes-+'
      +-number-+          '-No--'
      '-0-----'

>--+-----+-----+----->
+-EXPIration----yyyddd+
'-RETention----days----'

.-PROtection----No-----.-UNIT----3592-----.
```

```

>-----+-----+-----+-----+-----+-----+-----+-----+-----+-----><
'-PROtection---+No-----+-' '-UNIT---unit_name-'
      +-Yes-----+
      '-Automatic-'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the DEFINE LIBRARY command.

DEVType=3592 (Required)

Specifies the 3592 device type is assigned to the device class.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

See the following table for the recording formats.

Table 1. Recording formats for 3592

Format	Description
3592	Uncompressed (basic) format
3592C	Compressed format
3592-2	Uncompressed (basic) format, similar to the 3592 format
3592-C	Compressed format, similar to the 3592C format
3592-3	Uncompressed (basic) format, similar to the 3592 format
3592-3C	Compressed format, similar to the 3592C format
3592-4	Uncompressed (basic) format, similar to the 3592 format
3592-4C	Compressed format, similar to the 3592C format
DRIVE	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be different from the listed value.	

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. For optimal results, do not mix generations of drives in the same library. If a library contains mixed generations, media problems can result. For example, generation 1 and generation 2 drives cannot read generation 3 media. If possible, upgrade all drives to 3592 generation 3. If you cannot upgrade all drives to 3592 generation 3, you must use a special configuration.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is No. You can specify one of the following values:

Yes

Specifies that the drives use WORM media.

No

Specifies that the drives do not use WORM media.

Tip: The IBM Spectrum Protect™ server does not automatically delete scratch volumes in WORM storage pools after the volumes are emptied by expiration or other processes. To delete these volumes and remove them from WORM storage

pools, you must use the DELETE VOLUME command. IBM Spectrum Protect cannot reuse WORM volumes that were written to by the server and then deleted from a storage pool.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: **K** (KB), **M** (MB), **G** (GB), or **T** (TB). For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G. The smallest value that is accepted is 100 KB (ESTCAPACITY=100K).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMpression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as *2014007* (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3592 tape. This parameter is optional. The default value is 3592. The unit name can be up to 8 characters.

AIX | Linux

DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server)

To use a z/OS® media server to access StorageTek drives such as the StorageTek T9840 or T10000, you must define an ECARTRIDGE device class. In the device class definition, specify a library that was defined with the LIBTYPE=ZOSMEDIA parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine DEVclass--device_class_name----->
>>-LIBRary----zos_media_library--DEVType----ECARTridge----->
.-FORMAT----DRIVE----- .-ESTCAPacity----9G---
>--+-----+-----+-----+----->
'-FORMAT----+DRIVE----+' '-ESTCAPacity----size-'
      +-T9840C----+
      +-T9840C-C--+
      +-T9840D----+
      +-T9840D-C--+
      +-T10000A---+
      +-T10000A-C++
      +-T10000B---+
      +-T10000B-C++
      +-T10000C---+
      +-T10000C-C++
      +-T10000D---+
      +-T10000D-C-'

.-PREFIX----ADSM-----
>--+-----+-----+-----+----->
'-PREFIX----+ADSM-----+'
      '-tape_volume_prefix-'

.-MOUNTRetention----60----- .-MOUNTWait----60-----
>--+-----+-----+-----+----->
'-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'

.-MOUNTLimit----2----- .-COMPression----Yes-----
>--+-----+-----+-----+----->
```

```

'-MOUNTLimit-----+DRIVES-+-'  '-COMPrEsson-----+Yes-+-'
      +-number-+                '-No--'
      '-0-----'

>--+-----+----->
+-EXPIration-----yyyyddd+
'-RETention-----days-----'

.-PROtEction-----No----- . -UNIT-----9840----- .
>--+-----+----->>
'-PROtEction-----+No-----'  '-UNIT-----unit_name-'
      +-Yes-----+
      '-Automatic-'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the DEFINE LIBRARY command.

DEVType=ECARTridge (Required)

Specifies that the ECARTRIDGE device type is assigned to the device class. The ECARTRIDGE device type is for StorageTek drives such as the StorageTek T9840 or T10000.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. See the following table for the recording formats.

Table 1. Recording formats for ECARTRIDGE tapes

Format	Estimated Capacity	Description
DRIVE	-	The server selects the highest format that is supported by the drive on which a volume is mounted. DRIVE is the default value. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge

Format	Estimated Capacity	Description
<p>Note:</p> <ul style="list-style-type: none"> Some formats use a compression feature of the tape drive hardware. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value. T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats. T10000D drives can read, but cannot write, the T10000A, T10000B, and T10000C formats. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional. The default estimated capacity is 9 GB.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: **K** (KB), **M** (MB), **G** (GB), or **T** (TB). For example, specify that the estimated capacity is 9 GB with the parameter `ESTCAPACITY=9G`. The smallest value that is accepted is 100 KB (`ESTCAPACITY=100K`).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is `ADSM`. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

`AB.CD2.E`

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is `ADSM.BFS`.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (`LIBTYPE=EXTERNAL`), do not specify the `MOUNTWAIT` parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is YES.

You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as *2014007* (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the

server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support ECARTRIDGE tapes. Use the unit name that represents the subset of drives in the library that are attached to the z/OS system. This parameter is optional. The default value is 9840. The unit name can be up to 8 characters.

Example: Define a device class with the ECARTRIDGE device type

Define a device class named E1 with the ECARTRIDGE device type and with RACF protection active for all tape volumes that are assigned to this device class. All data is compressed for this device class. The device class is for a z/OS media server library named ZOSELIB.

```
define devclass e1 devtype=ecartridge library=zoselib compression=yes
    protection=yes
```

AIX

Linux

DEFINE DEVCLASS (Define a FILE device class for z/OS media server)

To use a z/OS® media server to access storage volumes on magnetic disk devices, you must define a FILE device class. In the device class definition, specify a library that was defined with the LIBTYPE=ZOSMEDIA parameter.

A volume in this device class is a Virtual Storage Access Method (VSAM) linear data set that is accessed by the z/OS media server. SCRATCH volumes can be used with device class and the z/OS media server can dynamically allocate the VSAM LDS. It is not necessary to define volumes for the server to use the device class. If you define volumes, set the high-level qualifier (HLQ) so that SMS recognizes the allocation request by the z/OS media server. If you are using defined volumes, the format volume function is not supported for the server when this device class is used. The z/OS media server uses a FormatWrite feature of DFSMS Media Manager when filling FILE volumes.

You can define volumes for the FILE device class by using the DEFINE VOLUME command. However, the z/OS media server does not allocate space for a defined volume until the volume is opened for its first use.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```

>>-DEfINE DEVclass--device_class_name--DEVType---FILE----->
                                     .-MAXCAPacity---10G--.
>--LIBRary---library_name--+-----+----->
                                     '-MAXCAPacity---size-'

    .-PRIMARYalloc---2600M-.    .-SECONDARYalloc---2600M-.
>--+-----+-----+----->
    '-PRIMARYalloc---size--'    '-SECONDARYalloc---size--'

    .-PREFIX---ADSM-----
>--+-----+----->
    '-PREFIX---file_volume_prefix-'

    .-MOUNTLimit---20-----
>--+-----+----->>
    '-MOUNTLimit---number-'

```

Parameters

DEVType=FILE (Required)

Specifies that the FILE device type is assigned to the device class.

LIBRARY (Required)

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The disk storage that is used by this device class is accessed by the z/OS media server and managed by SMS.

For information about defining a library, see the DEFINE LIBRARY command.

MAXCAPacity

Specifies the maximum size of file volumes that are defined to a storage pool in this device class. This parameter is optional. The default value is 10 GB (MAXCAPACITY=10G).

Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum size is 1 MB (MAXCAPACITY=1M). The maximum size is 16384 GB (MAXCAPACITY=16384G).

PRIMARYalloc

Specifies the initial amount of space that is dynamically allocated when a new volume is opened. Enough space must be available to satisfy the primary allocation amount. Storage Management Subsystem (SMS) policy determines whether multiple physical volumes can be used to satisfy the primary allocation request.

This parameter is optional. Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum size is 100 KB (PRIMARYALLOC=100K). The maximum size is 16384 GB (MAXCAPACITY=16384G). The default size is 2600 MB (PRIMARYALLOC=2600M). All values are rounded to the next higher multiple of 256 KB.

To avoid wasted space, the dynamic allocation operation uses the smaller of the values that are specified in the two parameters, PRIMARYALLOC and MAXCAPACITY.

SMS automatic class selection (ACS) routines can affect whether the PRIMARYALLOC and SECONDARYALLOC parameter values are used.

SECONDARYalloc

Specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up. The data set for a file volume is extended up to the size set by the MAXCAPACITY parameter, then the volume is marked full.

Because secondary allocation of a linear data set cannot span a physical volume, consider the size of the physical volume when you select a secondary allocation size. For example, physical volumes for a 3390 Model 3 are approximately 2.8 GB. To ensure that each extend request occupies nearly an entire physical volume but not more, use a secondary allocation size that is just less than 2.8 GB. A secondary allocation amount of 2600 MB allots enough space for the VSAM volume data set (VVDS), the volume label, and the volume table of contents (VTOC).

This parameter is optional. Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum value is 0 KB (SECONDARYALLOC=0K). The default value is 2600 MB. The maximum value is 16384 GB. Except for 0, all values are rounded to the next higher multiple of 256 KB.

If you specify 0 (SECONDARYALLOC=0), the file volume cannot be extended beyond the primary allocation amount.

SMS automatic class selection (ACS) routines can affect whether the PRIMARYALLOC and SECONDARYALLOC parameter values are used.

If you specify a value for the SECONDARYALLOCATION parameter that is not 0, or if you allow the value to default to 2600M, the SMS DATACLAS associated with the PREFIX identifier (for example, High Level Qualifier) must have the Extended Addressability (EA) attribute specified. Without the EA attribute, the SMS DATACLAS limits the allocation of the VSAM LDS FILE volume to the primary extent. (See the description of the PRIMARYALLOCATION parameter). With the data set limited to primary allocation size, the data set cannot be extended by the z/OS media server, and the volume is marked FULL before the maximum capacity is reached.

Restriction: Ensure that the values that you specify for the PRIMARYALLOC and SECONDARYALLOC parameters are within practical limits for the storage device. The server cannot check whether the values exceed practical device limits, and does not check whether the two values together exceed the current MAXCAPACITY setting.

Tip: To fill volumes when you specify a large value for the MAXCAPACITY parameter, specify large values for the PRIMARYALLOC and SECONDARYALLOC parameters. Use larger MVS™ volume sizes to reduce the chance of extend failure.

PREFIX

Specifies the high-level qualifier of the data set name that is used to allocate scratch volume data sets. For all scratch file volumes created in this device class, the server uses this prefix to create the data set name. This parameter is optional. The default is ADSM. The maximum length of the prefix, including periods, is 32 characters.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a file volume data set name using the default prefix is `ADSM.B0000021.BFS`.

If you have a data set naming convention, use a prefix that conforms to your naming conventions. For example, the following value is acceptable: `TSM.SERVER2.VSAMFILE`.

If you are running multiple server instances for either IBM Spectrum Protect™ or Tivoli® Storage Manager for z/OS Media you must use a unique value for the PREFIX parameter for each device class that you define.

MOUNTLimit

Specifies the maximum number of FILE volumes that can be open concurrently for this device class. This parameter is optional. The default value is 20.

If you are using IBM® 3995 devices that emulate 3390 devices, set the value no higher than the number of concurrent input or output streams that are possible on the physical media.

The value that you specify in this parameter is important if there is a significant penalty switching from one volume to another. For example, switching can take place when using IBM 3995 devices to emulate 3390 devices. The value that you specify must be no higher than the number of physical drives available on the device.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

DEFINE DOMAIN (Define a new policy domain)

Use this command to define a new policy domain. A policy domain contains policy sets, management classes, and copy groups. A client is assigned to one policy domain. The ACTIVE policy set in the policy domain determines the rules for clients that are assigned to the domain. The rules control the archive, backup, and space management services that are provided for the clients.

You must activate a policy set in the domain before clients assigned to the policy domain can back up, archive, or migrate files.

Privilege class

To issue this command, you must have system privilege.

contain an archive copy group. Also, specify that backup versions are retained for 60 days when management classes or copy groups are deleted and the default management class does not contain a backup copy group.

```
define domain prog1
description="Programming Group Domain"
backretention=60 archretention=90
```

Related commands

Table 1. Commands related to DEFINE DOMAIN

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.
QUERY DOMAIN	Displays information about policy domains.
UPDATE DOMAIN	Changes the attributes of a policy domain.

DEFINE DRIVE (Define a drive to a library)

Use this command to define a drive. Each drive is assigned to a library, and so the library must be defined before you issue this command.

A path must be defined after you issue the DEFINE DRIVE command to make the drive usable by IBM Spectrum Protect™. For more information, see DEFINE PATH (Define a path). If you are using a SCSI or VTL library type, see PERFORM LIBACTION (Define or delete all drives and paths for a library).

You can define more than one drive for a library by issuing the DEFINE DRIVE command for each drive. Stand-alone drives always require a manual library.

Windows Restriction: Before you issue the DEFINE DRIVE command, for a removable media device such as a Jaz, Zip, or CD drive, you must load the drive with properly formatted and labeled media.

For detailed and current drive support information, see the Supported Devices website for your operating system:

- **AIX** **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfIne DRive--library_name--drive_name----->
      .-SERial----AUTODetect----- .-ONLine----Yes-----
>--+-----+-----+-----+-----+-----+----->
      '-SERial----+AUTODetect----+' '-ONLine----+Yes--+'
              '-serial_number-'           '-No--'
                                (1)
      .-ELEMEnt----AUTODetect-----
>--+-----+-----+-----+-----+-----+----->
      '-ELEMEnt----+AUTODetect--+'
              '-address-----'
                                (2)
>--+-----+-----+-----+-----+-----+----->
      |                                     |
      '-ACSDRVID----drive_id-----'
```


Windows Restriction: To use ACSLS functions, the installation of StorageTek Library Attach software is required.

CLEANFREQUENCY

Specifies how often the server activates drive cleaning. This parameter is optional. For the most complete automation of cleaning for an automated library, you must have a cleaner cartridge that is checked into the library's volume inventory.

If you are using library-based cleaning, NONE is advised when your library type supports this function.

This parameter is not valid for externally managed libraries, such as 3494 libraries or StorageTek libraries that are managed under ACSLS.

Important: There are special considerations if you plan to use server-activated drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

NONE

Specifies that the server does not track cleaning for this drive. This value can be used for libraries that have their own automatic cleaning.

ASNEEDED

Specifies that the server loads the drive with a checked-in cleaner cartridge only when a drive reports to the device driver that it needs cleaning.

The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. See the Supported Devices website for your operating system to view detailed drive information. If ASNEEDED is not supported, you can use the gigabytes value for automatic cleaning.

For IBM 3592 and LTO drives, library-based cleaning is advised. If library-based cleaning is not supported, then ASNEEDED must be used. Gigabytes is not recommended.

Restriction: IBM Spectrum Protect does not control the drives that are connected to the NAS file server. If a drive is attached only to a NAS file server (no connection to a storage agent or server), do not specify ASNEEDED for the cleaning frequency.

gigabytes

Specifies, in gigabytes, how much data is processed on the drive before the server loads the drive with a cleaner cartridge. The server resets the gigabytes-processed counter each time it loads a cleaner cartridge in the drive. Important: When CLEANFREQUENCY=gigabyte, drive cleaning can occur before the gigabyte setting is reached, if the drive notifies the device driver that a cleaning is necessary.

Consult the information from the drive manufacturer for cleaning recommendations. If the information gives recommendations for cleaning frequency in terms of hours of use, convert to a gigabytes value by doing the following:

1. Use the bytes-per-second rating for the drive to determine a gigabytes-per-hour value.
2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

Using the cleaning frequency that is recommended by IBM® for IBM drives ensures that the drives are not overcleaned.

For IBM 3590 drives, specify a gigabyte value for the cleaning frequency to ensure that the drives receive adequate cleaning.

Example: Define a drive to library

Define a drive in a manual library with a library name of LIB01 and a drive name of DRIVE01.

```
define drive lib01 drive01
```

AIX

```
define path server01 drive01 srctype=server desttype=drive  
library=lib01 device=/dev/rmt0
```

Linux

```
define path server01 drive01 srctype=server desttype=drive  
library=lib01 device=/dev/tmscsi/mt0
```

Windows

```
define path server01 drive01 srctype=server desttype=drive
library=lib01 device=mt3.0.0.0
```

Example: Define a drive in an ACSLS library

Define a drive in an ACSLS library with a library name of ACSLIB and a drive name of ACSDRV1.

```
define drive acslib acsdrv1 acsdrv1=1,2,3,4
```

AIX

```
define path server01 acsdrv1 srctype=server desttype=drive
library=acslib device=/dev/rmt0
```

Linux

```
define path server01 acsdrv1 srctype=server desttype=drive
library=acslib device=/dev/tmscsi/mt0
```

Windows

```
define path server01 acsdrv1 srctype=server desttype=drive
library=acslib device=mt3.0.0.0
```

Example: Define a drive in an automated library

Define a drive in an automated library with a library name of AUTO8MMLIB and a drive name of DRIVE01.

```
define drive auto8mmlib drive01 element=82
```

AIX

```
define path server01 drive01 srctype=server desttype=drive
library=auto8mmlib device=/dev/rmt0
```

Linux

```
define path server01 drive01 srctype=server desttype=drive
library=auto8mmlib device=/dev/tmscsi/mt0
```

Windows

```
define path server01 drive01 srctype=server desttype=drive
library=auto8mmlib device=mt3.0.0.0
```

Related commands

Table 1. Commands related to DEFINE DRIVE

Command	Description
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE PATH	Changes the attributes associated with a path.

DEFINE EVENTSERVER (Define a server as the event server)

Use this command to identify a server as the event server.

If you define an event server, one IBM Spectrum Protect™ server can send events to another IBM Spectrum Protect server that will log those events.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine EVENTSERVer--server_name-----><
```

Parameters

server_name (Required)

Specifies the name of the event server. The server you specify must have already been defined with the DEFINE SERVER command.

Example: Designate the event server

Designate ASTRO to be the event server.

```
define eventserver astro
```

Related commands

Table 1. Commands related to DEFINE EVENTSERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE EVENTSERVER	Deletes reference to the event server.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
PING SERVER	Tests the connections between servers..
QUERY EVENTSERVER	Displays the name of the event server.
QUERY SERVER	Displays information about servers.

Related information:

[Enterprise event logging: logging events to another server](#)

DEFINE GRPMEMBER (Add a server to a server group)

Use this command to add a server as a member of a server group. You can also add one server group to another server group. A server group lets you route commands to multiple servers by specifying only the server group name.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine GRPMEMber--group_name---member_name+-----><
```

Parameters

group_name (Required)

Specifies the name of the server group to which the member will be added.

member_name (Required)

Specifies the names of the servers or groups to be added to the group. To specify multiple servers and groups, separate the names with commas and no intervening spaces. The servers or server groups must already be defined to the server.

Example: Define a server to a server group

Define the server SANJOSE to server group CALIFORNIA.

```
define grpmember california sanjose
```

Example: Define a server and a server group to a server group

Define the server TUCSON and the server group CALIFORNIA to server group WEST_COMPLEX.

```
define grpmember west_complex tucson,california
```

Related commands

Table 1. Commands related to DEFINE GRPMEMBER



Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DEFINE LIBRARY (Define a library)




Use this command to define a library. A library is a collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

A library can be accessed by only one source: an IBM Spectrum Protect™ server or a data mover. However, the drives in a library can be accessed by multiple sources.

The following library types can be defined to the server. Syntax and parameter descriptions are available for each type.

- DEFINE LIBRARY (Define a 349X library)
- DEFINE LIBRARY (Define an ACSLS library)
- DEFINE LIBRARY (Define an External library)
- DEFINE LIBRARY (Define a FILE library)
- DEFINE LIBRARY (Define a manual library)
- DEFINE LIBRARY (Define a SCSI library)
- DEFINE LIBRARY (Define a shared library)
- DEFINE LIBRARY (Define a VTL library)
-   DEFINE LIBRARY (Define a ZOSMEDIA library type)

For detailed and current library support information, see the Supported Devices website for your operating system:

-   Supported devices for AIX and Windows
-  Supported devices for Linux

To automatically label tape volumes in SCSI-type libraries, use the AUTOLABEL parameter on the DEFINE LIBRARY and UPDATE LIBRARY commands. Using this parameter eliminates the need to pre-label a set of tapes. It is also more efficient than using the LABEL LIBVOLUME command, which requires you to mount volumes separately. If you use the AUTOLABEL parameter, you must check in tapes by specifying CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

A label cannot include embedded blanks or periods and must be valid when used as a file name on the media.

You must label CD-ROM, Zip, or Jaz volumes with the device utilities from the manufacturer or the Windows utilities because IBM Spectrum Protect does not provide utilities to format or label these media types. The operating system utilities include the Disk Administrator program (a graphical user interface) and the label command.

Related commands

Table 1. Commands related to DEFINE LIBRARY

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE PATH	Changes the attributes associated with a path.

DEFINE LIBRARY (Define a 349X library)

Use this syntax to define a 349X library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType-----349X----->
      .-SHAREd-----No----- .-RESETDrives-----No-----.
```

```

>-----+-----+-----+-----+----->
'-SHAREd-----+Yes-+-' | (1) |
      '-No--'      '-RESETDrives-----+Yes-+-'
                                '-No--'

.-AUTOLabel-----Yes-----
>-----+-----+-----+-----+----->
'-AUTOLabel-----+No-----+-'
      +-Yes-----+
      '-OVERWRITE-'

.-SCRATCHCATegory-----301-----
>-----+-----+-----+-----+----->
'-SCRATCHCATegory-----number-'

.-PRIVATECATegory-----300-----
>-----+-----+-----+-----+----->
'-PRIVATECATegory-----number-'

>-----+-----+-----+-----+-----><
'-WORMSCRatchcategory-----number-'

```

Notes:

1. The default value of the RESETDRIVES parameter is conditional. If the SHARED parameter is set to NO, the value of the RESETDRIVES parameter is NO. If the SHARED parameter is set to YES, the value of the RESETDRIVES parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=349X (Required)

AIX | **Linux** Specifies that the library is an IBM 3494 or 3495 Tape Library Dataserver.

Windows Specifies that the library is an IBM 3494 Tape Library Dataserver or an IBM Tape System Library Manager emulating a 3494 Tape Library Dataserver.

Restriction: IBM 3494 libraries support only one unique device type at a time.

SHARED

Specifies whether this library is shared with other servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels only if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

SCRATCHCATegory

Specifies the category number to be used for scratch volumes in the library. This parameter is optional. The default value is 301 (becomes X'12D' on the IBM 3494 since it uses hexadecimal values). You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library.

PRIVATECATegory

Specifies the category number for private volumes that must be mounted by name. This parameter is optional. The default value is 300 (this value becomes X'12C' on the IBM 3494 because it uses hexadecimal values). You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library.

WORMSCRatchcategory

Specifies the category number to be used for WORM scratch volumes in the library. This parameter is required if you use WORM volumes. You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library. This parameter is only valid when 3592 WORM volumes are used.

Restriction: If the WORMSCRATCHCATEGORY is not defined and the WORM parameter is set to YES for the device class, the mount operation fails with an error message.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

AIX | Windows If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

The following table describes the three possible configurations for drives that are attached to NAS devices.

Table 1. Configurations for drives that are attached to NAS devices.

Library device configuration	The behavior for persistent reserve
The library device is attached to the IBM Spectrum Protect server, and the tape drives are shared by the server and the NAS device.	Drive reservation preemption is supported when the NAS device supports persistent reserve and it is enabled. For more information about setting persistent reserve, see the documentation for your NAS device.
The library device is attached to the IBM Spectrum Protect server and the tape drives are accessed only from the NAS device.	Drive reservation preemption is not supported. If you enable persistent reserve on the NAS device for these drives and a reservation is set by the NAS device but never cleared, you must use another method to clear the reservation.

AIX | Windows

Yes

Specifies that drive preemption through persistent reserve or target reset are used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve or target reset are not used. NO is the default for a library that is defined with SHARED=NO. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

Example: Define a 3494 library

Define a library named `my3494` with a scratch category number of 550, a private category number of 600, and a WORM scratch category number of 400®

```
define library my3494 libtype=349x scratchcategory=550
privatecategory=600 wormscratchcategory=400
```

DEFINE LIBRARY (Define an ACSLS library)

Use this syntax to define an ACSLS library.

Privilege class

Windows To use ACSLS functions, the installation of StorageTek Library Attach software is required.

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRary--library_name--LIBType----ACSLs----->
. -SHARed-----No----- . -RESEtDrives-----No-----
>--+-----+-----+-----+-----+----->
' -SHARed-----+Yes--+ ' | (1) |
      '-No--'      '-RESEtDrives-----+Yes--+-----'
                               '-No--'

. -AUTOLabel-----Yes-----
>--+-----+-----+-----+-----+-----><
' -AUTOLabel-----+No-----+ '
      +-Yes-----+
      '-OVERWRITE-'
```

Notes:

1. The default value of the RESETDRIVES parameter is conditional. If the SHARED parameter is set to NO, the value of the RESETDRIVES parameter is NO. If the SHARED parameter is set to YES, the value of the RESETDRIVES parameter is YES.

Parameters

`library_name` (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

`LIBType=ACSLs` (Required)

Specifies that the library is a StorageTek library that is controlled by StorageTek Automated Cartridge System Library Software (ACSLs).

`SHARed`

Specifies whether this library is shared with other servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

The following table describes the three possible configurations for drives that are attached to NAS devices.

Table 1. Configurations for drives that are attached to NAS devices.

Library device configuration	The behavior for persistent reserve
The library device is attached to the IBM Spectrum Protect server, and the tape drives are shared by the server and the NAS device.	Drive reservation preemption is supported when the NAS device supports persistent reserve and it is enabled. For more information about setting persistent reserve, see the documentation for your NAS device.
The library device is attached to the IBM Spectrum Protect server and the tape drives are accessed only from the NAS device.	Drive reservation preemption is not supported. If you enable persistent reserve on the NAS device for these drives and a reservation is set by the NAS device but never cleared, you must use another method to clear the reservation.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset are used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve or target reset are not used. NO is the default for a library that is defined with SHARED=NO. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

ACSID (Required)

Specifies the number of this StorageTek library that is assigned by the ACSSA (Automatic Cartridge System System Administrator). This number can be from 0 to 126. Issue QUERY ACS on your system to get the number for your library ID. This parameter is required.

For more information, see your StorageTek documentation.

Example: Define a shared ACSLS library

Define a library named ACSLIB with the library type of ACSLS and an ACSID of 1.

```
define library acslib libtype=acsls acsid=1 shared=yes
```

DEFINE LIBRARY (Define an External library)

Use this syntax to define an External library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType-----EXTernal----->
      .-AUTOLabel-----Yes-----
>---+-----+-----+-----+-----+-----+-----+----->>
      '-AUTOLabel-----+No-----+
                               +-Yes-----+
                               '-OVERWRITE-'
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=EXTernal (Required)

Specifies that the library is managed by an external media management system. This library type does not support drive definitions with the DEFINE DRIVE command. Rather, the external media management system identifies the appropriate drive for media access operations.

AIX | **Windows** In an IBM Spectrum Protect™ for Storage Area Networks environment, this parameter specifies that StorageTek Automated Cartridge System Library Software (ACSLS) or Library Station software controls the library. Software, such as Gresham EDT-DistribuTAPE, allows multiple servers to share the library. The drives in this library are not defined to IBM Spectrum Protect. ACSLS identifies the drive for media operations.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

Example: Define an external library for a SAN configuration

For an IBM Spectrum Protect for Storage Area Networks configuration, define a library named EXTLIB with the library type of EXTERNAL. If you are using Gresham Enterprise DistribuTAPE, the external library manager executable file is in the following directory:

- **AIX** /usr/lpp/dtelm/bin/elm
- **Linux** /opt/OMIdtelm/bin/elm
- **Windows** c:\program files\GES\EDT\bin\elm.exe

If you are using the IBM® Tape System Library Manager, the external library manager executable file can be found in the following directory:

- **AIX** **Linux** /opt/IBM/TSLM/client/tsm/elm
- **Windows** ...\\IBM\rmm\client\tsm\elm.exe

For more information, see the *IBM Tape System Library Manager User's Guide* at <http://www-01.ibm.com/support/docview.wss?uid=pub1ga32220802>.

1. Define the library:

```
define library extlib libtype=external
```

2. Define the path:

```
AIX  
define path server1 extlib srctype=server desttype=library  
externalmanager="/usr/lpp/dtelm/bin/elm"
```

```
Linux  
define path server1 extlib srctype=server desttype=library  
externalmanager="/opt/OMIdtelm/bin/elm"
```

```
Windows  
define path server1 extlib srctype=server desttype=library  
externalmanager="c:\program files\GES\EDT\bin\elm.exe"
```

DEFINE LIBRARY (Define a FILE library)

Use this syntax to define a FILE library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType---FILE----->  
.-SHAREd---No-----.  
>-----+-----><  
'-SHAREd---+Yes-+-'  
          '-No--'
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=FILE (Required)

Specifies that a pseudo-library is created for sequential file volumes. When you issue the DEFINE DEVCLASS command with DEVTYPE=FILE and SHARED=YES parameters, this occurs automatically. FILE libraries are necessary only when sharing sequential file volumes between the server and one or more storage agents. The use of FILE libraries requires

library sharing. Shared FILE libraries are supported for use in LAN-free backup configurations only. You cannot use a shared FILE library in an environment in which a library manager is used to manage library clients.

SHARED

Specifies whether this library is shared with other IBM Spectrum Protect™ servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

Example: Define a shared FILE library

Define a file library with shared=yes.

```
define library file1 libtype=file shared=yes
```

DEFINE LIBRARY (Define a manual library)

Use this syntax to define a manual library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType----MANUAL----->
. -RESETDrives----Yes-----
>--+-----+----->
' -RESETDrives----+Yes--+'
      '-No--'

. -AUTOLabel----Yes-----
>--+-----+-----><
' -AUTOLabel----+No-----+'
      +-Yes-----+
      '-OVERWRITE-'
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=MANUAL (Required)

Specifies that the library is not automated. When volumes must be mounted on drives in this type of library, messages are sent to operators. This type of library is used with stand-alone drives.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you need to check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset are used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve or target reset are not used. NO is the default for a library that is defined with SHARED=NO. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

Example: Define a manual library

Define a library named MANUALMOUNT with the library type of MANUAL.

```
define library manualmount libtype=manual
```

DEFINE LIBRARY (Define a SCSI library)

Use this syntax to define a SCSI library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFINE LIBRARY--library_name--LIBType-----SCSI----->
.-SHARED-----No----- .-RESETDrives-----No-----
>--+-----+-----+-----+-----+-----+-----+----->
  '-SHARED-----+Yes-+-' |                               (1) |
    '-No--'          '-RESETDrives-----+Yes-+-----'
                        '-No--'
```

```

.-AUTOLabel---No-----
>-----+----->
'-AUTOLabel---No-----'
          +-Yes-----+
          '-OVERWRITE-'

.-RELABELSCRatch---No-----
>-----+----->
'-RELABELSCRatch---No---+'
          '-Yes-'

.-SERial---AUTODetect-----
>-----+-----<<
'-SERial---AUTODetect---+'
          '-serial_number-'

```

Notes:

1. The default value of the RESETDRIVES parameter is conditional. If the SHARED parameter is set to NO, the value of the RESETDRIVES parameter is NO. If the SHARED parameter is set to YES, the value of the RESETDRIVES parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=SCSI (Required)

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, the server uses the media changer device.

SHARED

Specifies whether this library is shared with other servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is NO.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RELABELSCRatch

Specifies whether the server relabels volumes that were deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten. This parameter is optional and intended for use with a Virtual Tape Library (VTL).

If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might impact performance.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

No

Specifies that the server does not relabel volumes that are deleted and returned to scratch.

Yes

Specifies that the server relabels volumes that are deleted and returned to scratch.

RESETDrives

Specifies whether the server preempts a drive reservation if the drive is already reserved by persistent reserve when the server tries to access the drive. For example, a storage agent becomes unavailable, but the agent still holds the drive that is reserved through persistent reserve. With persistent reserve, the server can break a drive reservation and access the drive.

AIX | **Windows** If the drive is reserved by a SCSI-2 reserve, (and not by persistent reserve), the server uses a LUN reset to break the drive reservation to access the target device.

Linux LUN resets are not supported by the Linux operating system. If a drive is reserved by a SCSI-2 reserve, (and not by persistent reserve), the server is unable to break the reservation to access the drive. In this case, you can break the reservation by power cycling the device.

For network-attached storage (NAS) devices, reservation is controlled by the NAS file server. IBM Spectrum Protect™ does not control NAS devices and the RESETDrives parameter is not relevant for NAS devices.

Support for persistent reserve has the following limitations:

- If you are using the IBM Spectrum Protect device driver, persistent reserve is supported only on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. For information about driver configuration, see the *IBM Tape Device Drivers Installation and User's Guide*.
- If you are using a virtual tape library that is emulating a supported drive, persistent reserve might not be supported.
- A library manager is not able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reserve.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve or target reset is not used. NO is the default for a library that is defined with SHARED=NO. The RESETDrives parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

SERIAL

Specifies the serial number for the library that is being defined. This parameter is optional. The default is AUTODETECT.

If SERIAL=AUTODETECT, then when you define the path to the library, the serial number reported by the library is used as the serial number.

If SERIAL=*serial_number*, then the number you entered is compared to the number detected by the server.

Attention: Depending on the capabilities of the device, SERIAL=AUTODETECT might not be supported. In this case, the serial number is reported as blank.

Example: Define a SCSI library

Define a library that is named SCsilib with a library type of SCSI.

```
define library scsilib libtype=scsi
```

The library requires a path. The device name for the library is:

- **AIX** /dev/lb0
- **Linux** /dev/tmsmcsi/lb0
- **Windows** lb3.0.0.0

Define the path:

AIX

```
define path server1 scsilib srctype=server desttype=library
  device=/dev/lb0
```

Linux

```
define path server1 scsilib srctype=server desttype=library
  device=/dev/tmsmcsi/lb0
```

Windows

```
define path server1 scsilib srctype=server desttype=library
  device=lb3.0.0.0
```

DEFINE LIBRARY (Define a shared library)

Use this syntax to define a shared library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType-----SHARED----->
>>-PRIMarylibmanager-----server_name-----<<
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=SHARED (Required)

Specifies that the library is shared with another IBM Spectrum Protect™ server over a storage area network (SAN) or a dual SCSI connection to library drives.

Important: Specify this library type when you define the library on a library client.

PRIMarylibmanager

Specifies the name of the IBM Spectrum Protect server that is responsible for controlling access to library resources. You must define this server with the DEFINE SERVER command before you can use it as a library manager. This parameter is required and valid only if LIBTYPE=SHARED.

Example: Define a shared library

In a SAN, define a library named SHAREDTSM to a library client server named LIBMGR1

```
define library sharedtsm libtype=shared primarylibmanager=libmgr1
```

DEFINE LIBRARY (Define a VTL library)

Use this syntax to define a library that has a SCSI-controlled media changer device that is represented by a virtual tape library (VTL).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfINE LIBRARY--library_name--LIBType---VTL----->
. -SHARed---No----- . -RESEtDrives---No-----
>+-----+-----+-----+-----+----->
' -SHARed---+Yes-+' | (1) |
      '-No--'      '-RESEtDrives---+Yes-+'
                          '-No--'

. -AUTOLabel---No-----
>+-----+-----+-----+-----+----->
' -AUTOLabel---+No-----+'
      +-Yes-----+
      '-OVERWRITE-'

. -RELABELSCRatch---Yes-----
>+-----+-----+-----+-----+----->
' -RELABELSCRatch---+No-+-'
      '-Yes-'

. -SERial---AUTODetect-----
>+-----+-----+-----+-----+----->>
' -SERial---+AUTODetect-+-'
      '-serial_number-'
```

Notes:

1. The default value of the RESETDRIVES parameter is conditional. If the SHARED parameter is set to NO, the value of the RESETDRIVES parameter is NO. If the SHARED parameter is set to YES, the value of the RESETDRIVES parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=VTL (Required)

Specifies that the library has a SCSI-controlled media changer device that is represented by a virtual tape library. To mount volumes in drives in this type of library, the server uses the media changer device.

If you are defining a VTL library, your environment must not include any mixed-media and paths must be defined between all drives in the library and all defined servers, including storage agents, that use the library. If either of these characteristics are not true, the overall performance can degrade to the same levels as the SCSI library type; especially during times of high stress.

SHARed

Specifies whether this library is shared with other servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

RESEtDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

AIX | Windows

Yes

Specifies that drive preemption through persistent reserve or target reset are used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve or target reset are not used. NO is the default for a library that is defined with SHARED=NO. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES.

No

Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is NO.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RELABELSCRatch

Specifies whether the server relabels volumes that were deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten.

If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might impact performance.

Restriction: If you are defining a library that has drives that are attached to a network-attached storage (NAS) device, you must use the LABEL LIBVOLUME command to label the volumes for this library.

Yes

Specifies that the server relabels volumes that are deleted and returned to scratch. YES is the default.

No

Specifies that the server does not relabel volumes that are deleted and returned to scratch.

SERial

Specifies the serial number for the library that is being defined. This parameter is optional. The default is AUTODETECT.

If SERIAL=AUTODETECT, then when you define the path to the library, the serial number reported by the library is used as the serial number.

If SERIAL=*serial_number*, then the number you entered is compared to the number detected by the server.

Attention: Depending on the capabilities of the device, SERIAL=AUTODETECT might not be supported. In this case, the serial number is reported as blank.

Example: Define a VTL library

Define a library named VTL LIB with a library type of VTL.

```
define library vtllib libtype=vtl
```

The library requires a path. The device name for the library is:

- **AIX** /dev/lb0
- **Linux** /dev/tmsmcsi/lb0
- **Windows** lb3.0.0.0

Define the path:

AIX

```
define path server1 vtllib srctype=server desttype=library  
device=/dev/lb0
```

Linux

```
define path server1 vtllib srctype=server desttype=library  
device=/dev/tmsmcsi/lb0
```

Windows

```
define path server1 vtllib srctype=server desttype=library  
device=lb3.0.0.0
```

AIX

Linux

DEFINE LIBRARY (Define a ZOSMEDIA library type)

Use this syntax to define a library that represents a TAPE or FILE storage resource that is maintained by Tivoli® Storage Manager for z/OS® Media.

Define a library of type ZOSMEDIA when you want the library to be exclusively managed by Tivoli Storage Manager for z/OS Media. The library appears to the IBM Spectrum Protect™ server as a logical storage device that does not require DRIVE definitions. A PATH definition is required for the server and any storage agents that need access to the ZOSMEDIA library resource.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine LIBRARY--library_name--LIBType-----ZOSMEDIA-----><
```

Parameters

library_name (Required)

Specifies the name of the library to be defined.

LIBType=ZOSMEDIA (Required)

Specifies that the library type is the ZOSMEDIA which represents a TAPE or FILE storage resource that is maintained by Tivoli Storage Manager for z/OS Media.

Example: Configure a ZOSMEDIA library

The following example shows the steps needed to define and configure a zosmedia library. The configuration includes these components:

- A server named sahara
- A library defined as type zosmedia named zebra
- A z/OS media server named oasis
- A storage agent named mirage

Define a library named ZEBRA with a library type of ZOSMEDIA:

```
define library zebra libtype=zosmedia
```

Define the z/OS media server:

```
define server oasis serverpassword=sanddune
hladdress=9.289.19.67 lladdress=1777
```

The server requires a path to the library resource managed by Tivoli Storage Manager for z/OS Media:

```
define path sahara zebra srctype=server
desttype=library zosmediaserver=oasis
```

The storage agent requires a path to the library resource managed by Tivoli Storage Manager for z/OS Media:

```
define path mirage zebra srctype=server
desttype=library zosmediaserver=oasis
```

DEFINE MACHINE (Define machine information for disaster recovery)

Use this command to save disaster recovery information for a server or client node machine. This information will be included in the plan file to help you recover your machines.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine MACHine--machine_name----->
>--+-----+--+-----+--+-----+----->
  '-DESCRiption---description-' '-BUilding---building-'
>--+-----+--+-----+----->
  '-FLoor---floor-' '-ROom---room-'

.-PRIority---50----- .-ADSMServer---No-----
>--+-----+--+-----+-----<<
  '-PRIority---number---' '-ADSMServer---No---'
                                     '-Yes-'
```

Parameters

machine_name (Required)

Specifies the machine name. The name can be up to 64 characters.

DESCRiption

Specifies a machine description. This parameter is optional. The text can be up to 255 characters. Enclose the text in quotation marks if it contains any blank characters.

BUilding

Specifies the building that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

FLoor

Specifies the floor that this machine is on. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

ROom

Specifies the room that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRIority

Specifies the restore priority for the machine an integer from 1 to 99. The highest priority is 1. This parameter is optional. The default is 50.

ADSMServer

Specifies whether the machine is an IBM Spectrum Protect™ server. Only one machine can be defined as an IBM Spectrum Protect server. This parameter is optional. The default is NO. Possible values are:

No

This machine is not an IBM Spectrum Protect server.

Yes

This machine is an IBM Spectrum Protect server.

Example: Define a machine's disaster recovery information

Define a machine named DISTRICT5, and specify a location, a floor, and a room name. This machine contains critical data and has the highest priority.

```
define machine district5 building=101 floor=27
room=datafacilities priority=1
```

Related commands

Table 1. Commands related to DEFINE MACHINE

Command	Description
DEFINE MACHNODEASSOCIATION	Associates an IBM Spectrum Protect node with a machine.
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE MACHINE	Deletes a machine.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the IBM Spectrum Protect database.
QUERY MACHINE	Displays information about machines.
UPDATE MACHINE	Changes the information for a machine.

DEFINE MACHNODEASSOCIATION (Associate a node with a machine)

Use this command to associate client nodes with a machine. During disaster recovery, you can use this information to identify the client nodes that resided on destroyed machines.

The machine must be defined and the nodes registered to IBM Spectrum Protect™.

To retrieve the information, issue the QUERY MACHINE command. This information will be included in the plan file to help you recover the client machines.

A node remains associated with a machine unless the node, the machine, or the association itself is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
          .- - - - - .
          v           |
>>-DEFine MACHNODEAssociation--machine_name----node_name+----->>
```

Parameters

machine_name (Required)

Specifies the machine name.
node_name (Required)
Specifies the node names. A node can only be associated with one machine. To specify multiple nodes, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Example: Associate a node with a machine

Associate the node named ACCOUNTSPAYABLE with the machine named DISTRICT5.

```
define machnodeassociation district5 accountspayable
```

Related commands

Table 1. Commands related to DEFINE MACHNODEASSOCIATION

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DELETE MACHINE	Deletes a machine.
DELETE MACHNODEASSOCIATION	Deletes association between a machine and node.
QUERY MACHINE	Displays information about machines.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.

DEFINE MGMTCLASS (Define a management class)

Use this command to define a new management class in a policy set. To allow clients to use the new management class, you must activate the policy set that contains the new class.

You can define one or more management classes for each policy set in a policy domain. A management class can contain a backup copy group, an archive copy group, or both. The user of a client node can select any management class in the active policy set or use the default management class.

Attention: The DEFINE MGMTCLASS command fails if a copy storage pool is specified as the destination for files that were migrated by an IBM Spectrum Protect™ for Space Management client.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax

```
>>-DEFine MGmtclass--domain_name--policy_set_name--class_name--->
. -SPACEMGTEchnique----NONE----- .
>-----+----->
' -SPACEMGTEchnique----+AUTOMATIC+- '
      +-SElective+-
      ' -NONE----- '

. -AUTOMIGNOnuse----0---- .
>-----+----->
' -AUTOMIGNOnuse----days- '

. -MIGREQUIRESBkup----Yes---- .
>-----+----->
' -MIGREQUIRESBkup----+Yes+- '
      '-No-- '

```

```

.-MIGDESTination---SPACEMGPOOL-.
>---+-----+----->
'-MIGDESTination---pool_name---'
>---+-----+----->>
'-DESCRiption---description-'

```

Parameters

domain_name (Required)

Specifies the policy domain to which the management class belongs.

policy_set_name (Required)

Specifies the policy set to which the management class belongs. You cannot define a management class to the ACTIVE policy set.

class_name (Required)

Specifies the name of the new management class. The maximum length of this name is 30 characters. You cannot use either *default* or *grace_period* as a class name.

SPACEMGTECHnique

Specifies whether a file that is using this management class is eligible for migration. This parameter is optional. The default is NONE. This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients. Possible values are:

AUTOMATIC

Specifies that the file is eligible for both automatic migration and selective migration.

SElective

Specifies that the file is eligible for selective migration only.

NONE

Specifies that the file is not eligible for migration.

AUTOMIGNOnuse

Specifies the number of days that must elapse since a file was last accessed before it is eligible for automatic migration. This parameter is optional. The default value is 0. If SPACEMGTECHNIQUE is not AUTOMATIC, the server ignores this attribute. You can specify an integer in the range 0 - 9999.

This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients.

MIGREQUIRESBkup

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. The default is YES. This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients. Possible values are:

Yes

Specifies that a backup version must exist.

No

Specifies that a backup version is optional.

MIGDESTination

Specifies the primary storage pool where the server initially stores files that are migrated by IBM Spectrum Protect for Space Management clients. This parameter is effective only for IBM Spectrum Protect for Space Management clients, and is not effective for backup-archive clients or application clients. The default is SPACEMGPOOL.

Your choice for the destination might depend on factors such as the following:

- The number of client nodes that are migrated to the storage pool. When many user files are stored in the same storage pool, volume contention can occur as users try to migrate files to or recall files from the storage pool.
- How quickly the files must be recalled. If you need immediate access to migrated versions, you can specify a disk storage pool as the destination.

The command fails if you specify a copy storage pool or an active-data pool as the destination.

DESCRiption

Specifies a description of the management class. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a management class for a specific policy set and policy domain

Define a management class that is called MCLASS1 for policy set SUMMER in the PROG1 policy domain. For IBM Spectrum Protect for Space Management clients, allow both automatic and selective migration, and store migrated files in the SMPPOOL storage pool. Add the description, "Technical Support Mgmt Class."

```
define mgmtclass prog1 summer mclass1
spacemgtechnique=automatic migdestination=smpool
description="technical support mgmt class"
```

Related commands

Table 1. Commands related to DEFINE MGMTCLASS

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.

DEFINE NODEGROUP (Define a node group)

Use this command to define a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity. A node can be a member of one or more node groups.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
>>-DEFine NODEGroup--group_name----->
>--+-----+-----><
  '-DESCription----description-'
```

Parameters

group_name

Specifies the name of the node group that you want to create. The maximum length of the name is 64 characters. The specified name may not be the same as any existing client node name.

DESCription

Specifies a description of the node group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a node group

Define a node group named `group1`.

```
define nodegroup group1
```

Related commands

Table 1. Commands related to DEFINE NODEGROUP

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DEFINE NODEGROUPMEMBER (Define node group member)

Use this command to add a client node to a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity.

Privilege class

To issue this command you must have system or unrestricted policy privilege.

Syntax

```
          .-|-----|  
          v          |  
>>-DEFINE NODEGROUPMember--group_name----node_name+-----><
```

Parameters

`group_name`

Specifies the name of the node group to which you want to add a client node.

`node_name`

Specifies the name of the client node that you want to add to the node group. You can specify one or more names. Separate multiple names with commas; do not use intervening spaces. You can also use wildcard characters when specifying multiple names.

Example: Define node group members

Define two members, `node1` and `node2`, to a node group, `group1`.

```
define nodegroupmember group1 node1,node2
```

Related commands

Table 1. Commands related to DEFINE NODEGROUPMEMBER

Command	Description
---------	-------------

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DEFINE PATH (Define a path)

Use this command to define a path for a source to access a destination. Both the source and destination must be defined before you can define a path. For example, if a path is required between a server and a drive, you must first issue the DEFINE DRIVE command and then issue the DEFINE PATH command. A path must be defined after you issue the DEFINE DRIVE command in order to make the drive usable by the server.

Syntax and parameter descriptions are available for the following path types.

- DEFINE PATH (Define a path when the destination is a drive)
- DEFINE PATH (Define a path when the destination is a library)
- **AIX** | **Linux** DEFINE PATH (Define a path when the destination is a ZOSMEDIA library)

For detailed and current device support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Related commands

Table 1. Commands related to DEFINE PATH

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DATAMOVER	Changes the definition for a data mover.
UPDATE PATH	Changes the attributes associated with a path.

DEFINE PATH (Define a path when the destination is a drive)

Use this syntax when you define a path to a drive.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine PATH--source_name--destination_name----->
>--SRCType-----+DATAMover-+-----+----->
          '-SERVer----'   '-AUTODetect-----+No--+-'
                               '-Yes-'
>--DESTType-----DRive--LIBRARY-----library_name----->
>----DEVIce-----+device_name-+----->
          '-FILE-----'
          .-GENERICTAPE-----No----- .-ONLine-----Yes-----
>--+-----+-----+-----+----->
          '-GENERICTAPE-----+Yes-+-'   '-ONLine-----+Yes-+-'
                               '-No--'   '-No--'
          .-DIRectory-----current_directory_name-.
>--+-----+-----+-----><
          |           .-,-----|
          |           v           |
          '-DIRectory-----directory_name-+-----'
```

Parameters

source_name (Required)

Specifies the name of source for the path. This parameter is required.

destination_name (Required)

Specifies the name of the destination. This parameter is required.

SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVer

Specifies that a storage agent is the source.

AUTODetect

Specifies whether the serial number for a drive is automatically updated in the database at the time that the path is defined. This parameter is optional. This parameter is only valid for paths that are defined from the local server to a drive. Possible values are:

No

Specifies that the serial number is not automatically updated. The serial number is still compared with what is already in the database for the device. The server issues a message if there is a mismatch.

Yes

Specifies that the serial number is not automatically updated to reflect the same serial number that the drive reports to the server.

Important:

1. If you did not set the serial number when you defined the drive, the server always tries to detect the serial number, and AUTODETECT defaults to YES. If you previously entered a serial number, then AUTODETECT defaults to NO.
2. The use of AUTODETECT=YES in this command means that the serial number set in the drive definition is updated with the detected serial number.
3. If you set DESTTYPE=DRIVE and AUTODETECT=YES, then the drive element number in the database is automatically changed to reflect the same element number that corresponds to the serial number of that drive. This is true for drives in a SCSI library. For more information about the element number, see DEFINE DRIVE.
4. Depending on the capabilities of the device, the AUTODETECT parameter might not be supported.

DESTType=DRive (Required)

Specifies that a drive is the destination. When the destination is a drive, you must specify a library name.

LIBRARY

Specifies the name of the library to which the drive is assigned. The library and its drives must already be defined to the server. If the path is from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349X, or ACSLS.

DEVICE

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

AIX | **Windows** The source uses the device name to access the drive. See Table 1 for examples.

Table 1. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	AIX /dev/mt3 Windows mt3
Storage agent (on a Windows system) to a drive (not a FILE drive)	mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM® System Storage® N Series: rst01

Linux The source uses the device name to access the drive. See Table 2 for examples.

Table 2. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	/dev/tmscsi/mt3
Storage agent to a drive (not a FILE drive)	/dev/tmscsi/mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM System Storage N Series: rst01

Important:

- **AIX** | **Linux** For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. **Windows** For 349X libraries, the alias name is a symbolic name that is specified in the c:\winnt\ibmatl.conf file. For more information, see *IBM Tape Device Drivers Installation and User's Guide*, which can be downloaded from the IBM Systems support site at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the SYSCONFIG command. Use this command to determine device names for drives:

```
sysconfig -t
```

Windows GENERICTAPE

Windows Specifies whether the tape drive to be used is a GENERICTAPE device class type. If the device is a tape drive and is not supported by IBM Spectrum Protect™ but is supported for the Windows operating system, you can use it with the generic tape format. To use the drive, specify GENERICTAPE=Yes when you define a path to the drive. The default is No. Possible values are:

Yes

Specifies that the tape drive to be used is a GENERICTAPE device class type.

No

Specifies that the tape drive to be used is not a GENERICTAPE device class type.

ONLINE

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

For example, if the path from a data mover to a drive is online, but either the data mover or the drive is offline, you cannot use the path.

DIRectory

Specifies the directory location or locations where the storage agent reads and writes the files that represent storage volumes for the FILE device class that is associated with the FILE library. The DIRECTORY parameter is also used for devices of type REMOVABLEFILE. For REMOVABLEFILE devices, the DIRECTORY parameter provides information for the server (not a storage agent) along with the DRIVE parameter to describe access to the device. This parameter is optional. For a path from a storage agent to a FILE device, this parameter is only valid when *all* of the following conditions are true:

- The source type is SERVER (meaning a storage agent that has been defined as a server to this server).
- The source name is the name of a storage agent, *not* the server.
- The destination is a logical drive that is part of a FILE library that is created when the device class was defined.

If you specified multiple directories for the device class associated with the FILE library, you must specify the same number of directories for each path to the FILE library. Do not change or move existing directories on the server that the storage agent is using so that the device class and the path remain synchronized. Adding directories is permitted. Specifying a mismatched number of directories can cause a runtime failure.

The default value for DIRECTORY is the directory of the server at the time the command is issued. The Windows registry is used to locate the default value.

Use a naming convention that you can use to associate the directory with a particular physical drive. This can help ensure that your configuration is valid for sharing the FILE library between the server and storage agent. If the storage agent is on a Windows system, use a universal naming convention (UNC) name. When the storage agent lacks permission to access remote storage, it experiences mount failures.

Windows The account that is associated with the storage agent service must either be an account within the local administrator's group or an account within the domain administrator's group. If the account is in the local administrator's group, the user ID and password must match that of an account with permissions to access storage as provided by the system that administers the remote share. For example, if a SAMBA server is providing access to remote storage, the user ID and password in the SAMBA configuration must match that of the local administrator user ID and password associated with the storage agent service.

```
define devclass file devtype=file shared=yes mountlimit=1
directory=d:\filedir\dir1
define path stal file1 srctype=server desttype=drive
library=file1 device=file
directory=\\192.168.1.10\filedir\dir1
```

In the previous example, the DEFINE DEVCLASS command establishes the shared file system in the directory that is accessed by the server as D:\FILEDIR\DIR1. The storage agent, however, is using UNC name \\192.168.1.10\FILEDIR\DIR1. This means that the system with TCP/IP address 192.168.1.10 is sharing the same directory using FILEDIR as the shared name. Also, the storage agent service has an account that can access this storage. It can access it either because it is associated with a local account with the same user ID and password as 192.168.1.10 or it is associated with a domain account that is available on both the storage agent and on 192.168.1.10. If appropriate to the installation, you can replace the 192.168.1.10 with a symbolic name such as:

```
example.yourcompany.com
```

Attention:

1. Storage agents access FILE volumes by replacing a directory name in a volume name with a directory name from a directory in the list provided with the DEFINE PATH command. Directories that are specified with this parameter are not validated on the server.
2. IBM Spectrum Protect does not create shares or permissions, or mount the target file system. You must complete these actions before you start the storage agent.

Example: Define a path from a server to a drive

Define a path from a server to a drive. In this case, the server name is *NET1*, the drive name is *TAPEDRV6*, the library is *NETLIB*, and the device name is *mt4*. Set *AUTODETECT* to *NO*.

```
define path net1 tapedrv6 srctype=server autodetect=no desttype=drive
  library=netlib device=mt4
```

Example: Define a path from a data mover server to a drive for backup and restore

Define a path from the data mover that is a NAS file server to the drive that the NAS file server will use for backup and restore operations. In this example, the NAS data mover is *NAS1*, the drive name is *TAPEDRV3*, the library is *NASLIB*, and the device name for the drive is *rst0l*.

```
define path nas1 tapedrv3 srctype=datamover desttype=drive library=naslib
  device=rst0l
```

Linux

Example: Define a path from a storage agent to a drive for backup and restore

Define a path from storage agent *SA1* to the drive that the storage agent uses for backup and restore operations. In this example, the library is *TSMLIB*, the drive is *TAPEDRV4*, and the device name for the drive is */dev/tmscsi/mt3*.

```
define path sa1 tapedrv4 srctype=server desttype=drive library=tsmlib
  device=/dev/tmscsi/mt3
```

AIX | Windows

Example: Define a path from a storage agent to a drive for backup and restore

Define a path from storage agent *SA1* to the drive that the storage agent uses for backup and restore operations. In this example, the library is *TSMLIB*, the drive is *TAPEDRV4*, and the device name for the drive is */dev/mt3*.

```
define path sa1 tapedrv4 srctype=server desttype=drive library=tsmlib
  device=/dev/mt3
```

AIX | Windows

Example: Define a path to give a storage agent access to shared disk storage

Define a path that gives the storage agent access to files on disk storage that is shared with the server. Drive *FILE9* is defined to library *FILE1* on the server. The storage agent *SA1* accesses *FILE9*. On the storage agent, this data is on directory *\\192.168.1.10\filedata*.

AIX The data for *FILE9* resides on the server at */tsmdata/filedata*.

Windows The data for *FILE9* resides on the server at *d:\tsmdata\filedata*.

```
define path sa1 file9 srctype=server desttype=drive library=file1 device=file
  directory="\\192.168.1.10\filedata"
```

Example: Configure a storage agent to use a FILE library

The following example illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created *FILE* volumes.

Suppose you want to use these three directories for a *FILE* library: Windows

- c:\server
- d:\server
- e:\server

AIX | Linux

- /opt/tivoli1
- /opt/tivoli2
- /opt/tivoli3

1. Use the following command to set up a *FILE* library named *CLASSA* with one drive named *CLASSA1* on *SERVER1*: Windows

```
define devclass classa devtype=file
directory="c:\server,d:\server,e:\server"
shared=yes mountlimit=1
```

AIX Linux

```
define devclass classa devtype=file
directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
shared=yes mountlimit=1
```

2. You want the storage agent STA1 to be able to use the FILE library, so you define the following path for storage agent STA1:

Windows

```
define path sta1 classal srctype=server desttype=drive device=file
directory="\\192.168.1.10\c\server,\\192.168.1.10\d\server,
\\192.168.1.10\e\server" library=classa
```

AIX Linux

```
define path sta1 classal srctype=server desttype=drive device=file
directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

Windows In this scenario, the storage agent, STA1, replaces the directory name c:\server with the directory name \\192.168.1.10\c\server to access FILE volumes that are in the c:\server directory on the server.

AIX Linux

In this scenario, the storage agent, STA1, replaces the directory name /opt/tivoli1 with the directory name /opt/ibm1/ to access FILE volumes that are in the /opt/tivoli1 directory on the server.

3. **Windows** File volume c:\server\file1.dsm is created by SERVER1. If you later change the first directory for the device class with the following command:

```
update devclass classa directory="c:\otherdir,d:\server,e:\server"
```

SERVER1 is still able to access file volume c:\server\file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list that is associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

4. If file volume /opt/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,
/opt/tivoli3"
```

SERVER1 is still able to access file volume /opt/tivoli1/file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list that is associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

DEFINE PATH (Define a path when the destination is a library)

Use this syntax when defining a path to a library.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEFine PATH--source_name--destination_name----->
(1)
>--SRCType-----+--DATAMover-----+-----+----->
'-SERVer-----' '-AUTODetect-----+No---+'
'-Yes-'
>--DESTType-----LIBRARY---+--DEVIce-----+----->
'-EXTERNALManager---+--path_name-'
```


Source to destination	Example
NAS data mover to a library	mc0

Linux The source uses the device name to access the library. See Table 2 for examples.

Table 2. Examples of device names

Source to destination	Example
Server to a library	/dev/tsm SCSI/lb4
NAS data mover to a library	mc0

Important:

- **AIX** | **Linux** For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. **Windows** For 349X libraries, the alias name is a symbolic name that is specified in the c:\winnt\ibmatl.conf file. For more information, see *IBM Tape Device Drivers Installation and User's Guide*, which can be downloaded from the IBM® Systems support site at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the SYSCONFIG command. Use this command to determine device names for drives:

```
sysconfig -t
```

Use this command to determine the device name for a library:

```
sysconfig -m
```

EXTERNALManager

Specifies the location of the external library manager where IBM Spectrum Protect can send media access requests. Use single quotation marks around the value of this parameter. For example, enter: **AIX**

```
/usr/lpp/GESEdt-acsls/bin/elmdt
```

Linux

```
/opt/GESEdt-acsls/bin/elmdt
```

Windows

```
C:\Program Files\GES\EDT-ACSLs\bin\elmdt.exe
```

This parameter is required when the library name is an external library.

ONLine

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Attention: If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized.

Example: Define a path from a server to a library

Define a path from the server SATURN to the SCSI type library SCSILIB: **AIX**

```
define path saturn scsilib srctype=server
desttype=library device=/dev/lb3
```

Linux

```
define path saturn scsilib srctype=server
desttype=library device=/dev/tsm SCSI/lb3
```

Windows

```
define path saturn scsilib srctype=server
desttype=library device=lb3.0.0.0
```

AIX Linux

DEFINE PATH (Define a path when the destination is a ZOSMEDIA library)

Use this syntax when defining a path to a ZOSMEDIA library. You must first define the z/OS® media server in your configuration with the DEFINE SERVER command.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DEfInE PATH--source_name--destination_name----->
>--SRCType-----SERVer--DESTType-----LIBRary----->
>--ZOSMEDIASERVER-----server_name-----ONLine-----<
                                     .-ONLine-----Yes-----
                                     +-----+-----+-----+-----<
                                     '-ONLine-----+-Yes+-'
                                     '-No--'
```

Parameters

source_name (Required)

Specifies the name of source for the path.

destination_name (Required)

Specifies the name of the ZOSMEDIA library.

SRCType=SERVer (Required)

Specifies that a storage agent or server is the source.

DESTType=LIBRary (Required)

Specifies that a library is the destination.

ZOSMEDIAServer (Required)

Specifies the name of the server that represents a Tivoli® Storage Manager for z/OS Media server.

ONLine

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Attention: If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized.

If the z/OS media server cannot be accessed during initialization of the IBM Spectrum Protect™ server, the library path will be set offline. Use the UPDATE PATH command and specify ONLINE=YES to vary the ZOSMEDIA library back online.

DEFINE POLICYSET (Define a policy set)

Use this command to define a policy set in a policy domain. A policy set contains management classes, which contain copy groups. You can define one or more policy sets for each policy domain.

To put a policy set into effect, you must activate the policy set by using the ACTIVATE POLICYSET command. Only one policy set can be active in a policy domain. The copy groups and management classes within the active policy set determine the rules by which client nodes perform backup, archive, and space management operations, and how the client files stored are managed.

Use the VALIDATE POLICYSET command to verify that a policy set is complete and valid before activating it with the ACTIVATE POLICYSET command.

Privilege class

To issue this command you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-DEFine Policyset--domain_name--policy_set_name----->
>--+-----+-----><
  '-DESCription-----description-'
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the policy set belongs.

policy_set_name (Required)

Specifies the name of the policy set. The maximum length of this name is 30 characters. You cannot define a policy set named ACTIVE.

DESCription

Specifies a description for the new policy set. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a policy set

Define a policy set called `SUMMER` for the `PROG1` policy domain and include the description, "Programming Group Policies."

```
define policyset prog1 summer
description="Programming Group Policies"
```

Related commands

Table 1. Commands related to DEFINE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE MGMTCLASS	Defines a management class.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

DEFINE PROFASSOCIATION (Define a profile association)

Use this command on a configuration manager to associate one or more objects with a configuration profile for distribution to subscribing managed servers. After a managed server subscribes to a profile, the configuration manager sends object definitions associated with the profile to the managed server where they are stored in the database. Objects created this way in the database of a managed server become managed objects. An object can be associated with more than one profile.

You can use this command to define an initial set of profile associations and to add to existing associations.

You can associate the following types of objects with a profile:

- Administrator registrations and authorities
- Policy domains, which include the domains' policy sets, management classes, copy groups, and client schedules
- Administrative schedules
- Server command scripts
- Client option sets
- Server definitions
- Server group definitions

Tip: The configuration manager does not distribute status information for an object to managed servers. For example, information such as the number of days since an administrator last accessed the server is not distributed to managed servers. This type of information is maintained in the databases of the individual managed servers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine PROFASSOCIation--profile_name----->
>--+-----+----->
  '-ADMins---+*-----+'
      | .-,----- . |
      | V           | |
      '---admin_name+--'

>--+-----+----->
  '-DOMains---+*-----+'
      | .-,----- . |
      | V           | |
      '---domain_name+--'

>--+-----+----->
  '-ADSCHeds---+*-----+'
      | .-,----- . |
      | V           | |
      '---schedule_name+--'

>--+-----+----->
  '-SCRipts---+*-----+'
      | .-,----- . |
      | V           | |
      '---script_name+--'

>--+-----+----->
  '-CLOptsets---+*-----+'
      | .-,----- . |
      | V           | |
      '---option_set_name+--'

>--+-----+----->
  '-SERVers---+*-----+'
      | .-,----- . |
      | V           | |
      '---server_name+--'

>--+-----+-----><
  '-SERVERGroups---+*-----+'
      | .-,----- . |
      | V           | |
      '---group_name+--'
```

Parameters

profile_name (Required)

Specifies the name of the configuration profile.

ADMins

Specifies administrators to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all administrators that are registered with the configuration manager. If you specify the match-all definition and later add more administrators, they are automatically distributed through the profile.

The configuration manager distributes the administrator name, password, contact information, and authorities of administrators associated with the profile. The configuration manager does not distribute the following:

- The administrator named SERVER_CONSOLE, even if you use a match-all definition
- The locked or unlocked status of an administrator

When the profile already has administrators associated with it, the following apply:

- If you specify a list of administrators and a list already exists, IBM Spectrum Protect™ combines the new list with the existing list.
- If you specify a match-all definition and a list of administrators already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of administrators, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the ADMINS=* parameter.

DOmains

Specifies policy domains to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all domains that are defined on the configuration manager. If you specify the match-all definition and later add more domains, they are automatically distributed through the profile.

The configuration manager distributes domain information that includes definitions of policy domains, policy sets, management classes, copy groups, and client schedules. The configuration manager does not distribute the ACTIVE policy set. Administrators on a managed server can activate any policy set within a managed domain on a managed server.

When the profile already has domains associated with it, the following apply:

- If you specify a list of domains and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of domains already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of domains, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the DOMAINS=* parameter.

Important: Client operations such as backup and archive fail if destination pools do not exist. Therefore, managed servers that subscribe to this profile must have definitions for any storage pools specified as destinations in the associated domains. Use the RENAME STGPOOL command to rename existing storage pools to match the destination names distributed.

ADSCHeds

Specifies administrative schedules to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all administrative schedules that are defined on the configuration manager. If you specify the match-all definition and later add more administrative schedules, they are automatically distributed through the profile.

Tip: Administrative schedules are not active when they are distributed by a configuration manager. An administrator on a managed server must activate any schedule to have it run on that server.

When the profile already has administrative schedules associated with it, the following apply:

- If you specify a list of administrative schedules and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of administrative schedules already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of administrative schedules, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the ADSCHEDS=* parameter.

SCRipts

Specifies server command scripts to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all scripts that are defined on the configuration manager. If you specify the match-all definition and later add more scripts, they are automatically distributed through the profile.

When the profile already has scripts associated with it, the following apply:

- If you specify a list of scripts and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of scripts already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of scripts, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the SCRIPTS=* parameter.

CLOptsets

Specifies client option sets to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all client option sets that are defined on the configuration manager. If you specify the match-all definition and later add more client option sets, they are automatically distributed through the profile.

When the profile already has client option sets associated with it, the following apply:

- If you specify a list of client option sets and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of client option sets already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of client option sets, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the CLOPSETS=* parameter.

SERVers

Specifies server definitions to associate with the profile. The definitions are distributed to managed servers that subscribe to this profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all servers that are defined on the configuration manager. If you specify the match-all definition and later add more servers, they are automatically distributed through the profile.

The configuration manager distributes the following server attributes: communication method, IP address, port address, server password, URL, and the description. Distributed server definitions always have the ALLOWREPLACE attribute set to YES on the managed server, regardless of this parameter's value on the configuration manager. On the managed server, you can use the UPDATE SERVER command to set all other attributes.

When the profile already has servers associated with it, the following apply:

- If you specify a list of servers and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of servers already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of servers, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the SERVERS=* parameter.

Important:

1. A server definition on a managed server is not replaced by a definition from the configuration manager unless you have allowed replacement of the definition on the managed server. To allow replacement, on the managed server update the server definition by using the UPDATE SERVER command with ALLOWREPLACE=YES.
2. If a configuration manager distributes a server definition to a managed server, and a server group of the same name exists on the managed server, the distributed server definition replaces the server group definition.

SERVERGroups

Specifies server groups to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk

(*) by itself, to specify all server groups that are defined on the configuration manager. If you specify the match-all definition and later add more server groups, they are automatically distributed through the profile.

Tip: A configuration manager does not distribute a server group definition to a managed server if the managed server has a server defined with the same name as that of the server group.

When the profile already has server groups associated with it, the following apply:

- If you specify a list of server groups and a list already exists, IBM Spectrum Protect combines the new list with the existing list.
- If you use a match-all definition and a list of server groups already exists, IBM Spectrum Protect replaces the list with the match-all definition.
- If you specify a list of server groups, and a match-all definition had previously been specified, IBM Spectrum Protect ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the SERVERGROUPS=* parameter.

Example: Associate a specific domain with a specific profile

Associate a domain named MARKETING with a profile named DELTA.

```
define profassociation delta domains=marketing
```

Example: Associate all domains with a specific profile

You have already associated a list of domains with a profile named GAMMA. Now associate all domains defined on the configuration manager with the profile.

```
define profassociation gamma domains=*
```

Related commands

Table 1. Commands related to DEFINE PROFASSOCIATION

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DEFINE PROFILE (Define a profile)

Use this command on a configuration manager to define a profile (a set of configuration information) that can be distributed to managed servers.

After defining a profile, you can use the DEFINE PROFASSOCIATION command to specify objects to be distributed to managed servers subscribing to the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine PROFILE--profile_name----->
>--+-----+-----><
  '-DESCRiption----description-'
```

Parameters

profile_name (Required)

Specifies the name of the profile. The maximum length of the name is 30 characters.

DESCRiption

Specifies a description of the profile. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. This parameter is optional.

Example: Define a new profile

Define a profile named ALPHA with a description of "Programming Center."

```
define profile alpha
description="Programming Center"
```

Related commands

Table 1. Commands related to DEFINE PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)

Use this command to associate recovery media with one or more machines. A machine is associated with recovery media so that the location of the boot media and its list of volume names are available to recover the machine. To retrieve the information, issue the QUERY MACHINE command. This information will be included in the plan file to help you recover the client machines.

To associate a machine with recovery media, both the machine and media must be defined to IBM Spectrum Protect™. A machine remains associated with the media until the association, the media, or the machine is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
.------.
```

```
>>-DEFine RECMEDMACHAssociation--media_name----machine_name+--><
```

Parameters

media_name (Required)

Specifies the name of the recovery media with which one or more machines will be associated.

machine_name (Required)

Specifies the name of the machines to be associated with the recovery media. A machine can be associated with multiple recovery media. To specify a list of machines, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Example: Associate machines to recovery media

Associate machines DISTRICT1 and DISTRICT5 to the DIST5RM recovery media.

```
define recmedmachassociation dist5rm  
district1,district5
```

Related commands

Table 1. Commands related to DEFINE RECMEDMACHASSOCIATION

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
DELETE MACHINE	Deletes a machine.
DELETE RECMEDMACHASSOCIATION	Deletes association between recovery media and a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
QUERY MACHINE	Displays information about machines.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.

DEFINE RECOVERYMEDIA (Define recovery media)

Use this command to define the media needed to recover a machine. The same media can be associated with multiple machines. To display the information, use the QUERY MACHINE command. This information will be included in the plan file to help you to recover the client machines.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine RECOVERYMedia--media_name----->  
  
>--+-----+----->  
|           .-,-----.|  
|           v           ||  
'-VOLumename-----volume_name+--'  
  
>--+-----+-----+----->  
'-DEScRiption-----description-' '-LOcation-----location-'  
  
. -Type-----Other-----.  
>--+-----+-----+----->  
'-Type-----+Other+-' '-PRoduct-----product_name-'  
'-BOot--'  
  
>--+-----+-----><
```

'-PRODUCTInfo-----product_information-'

Parameters

media_name (Required)

Specifies the name of the recovery media to be defined. The name can be up to 30 characters.

VOLumenames

Specifies the names of volumes that contain the recoverable data (for example, operating system image copies). This parameter is required if you specify a media type of BOOT. Specify boot media volume names in the order in which they are to be inserted into the machine at recovery time. The maximum length of the volume names list is 255 characters. Enclose the list in quotation marks if it contains any blank characters.

DESCription

Specifies the description of the recovery media. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

LOcation

Specifies the location of the recovery media. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

Type

Specifies the type of recovery media. This parameter is optional. The default is OTHER.

BOot

Specifies that this is boot media. You must specify volume names if the type is BOOT.

OTHer

Specifies that this is not boot media. For example, a CD that contains operating system manuals.

PROduct

Specifies the name of the product that wrote to this media. This parameter is optional. The maximum length is 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRODUCTInfo

Specifies information about the product that wrote to the media. This would be information that you may need to restore the machine. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

Example: Define the media needed to recover a machine

Define the recovery media named DIST5RM. Include a description and the location.

```
define recoverymedia dist5rm
description="district 5 base system image"
location="district 1 vault"
```

Related commands

Table 1. Commands related to DEFINE RECOVERYMEDIA

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

DEFINE SCHEDULE (Define a client or an administrative command schedule)

Use this command to create a client or administrative command schedule.

The DEFINE SCHEDULE command takes two forms: one if the schedule applies to client operations, one if the schedule applies to administrative commands. Within these two forms, you can select either classic or enhanced style schedules. The syntax and

parameters for each form are defined separately.

For each schedule, a startup window is specified. The startup window is the time period during which the schedule must be initiated. The schedule will not necessarily complete processing within this window. If the server is not running when this window starts, but is started before the end of the defined window is reached, the schedule will run when the server is restarted. Options associated with each schedule style (classic and enhanced) determine when the startup windows should begin.

Table 1. Commands related to DEFINE SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE ASSOCIATION	Associates clients with a schedule.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY SCHEDULE	Displays information about schedules.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.
SET MAXSCHEDSESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.
UPDATE SCHEDULE	Changes the attributes of a schedule.

- **DEFINE SCHEDULE (Define a client schedule)**
Use the DEFINE SCHEDULE command to define a client schedule. IBM Spectrum Protect uses this schedule to automatically perform a variety of client operations for your client workstation at specified intervals or days. After you define a schedule, use the DEFINE ASSOCIATION command to associate the client with the schedule.
- **DEFINE SCHEDULE (Define a schedule for an administrative command)**
Use the DEFINE SCHEDULE command to create a new schedule for processing an administrative command.

DEFINE SCHEDULE (Define a client schedule)

Use the DEFINE SCHEDULE command to define a client schedule. IBM Spectrum Protect™ uses this schedule to automatically perform a variety of client operations for your client workstation at specified intervals or days. After you define a schedule, use the DEFINE ASSOCIATION command to associate the client with the schedule.

You must start the client scheduler on the client workstation for IBM Spectrum Protect to process the schedule.

Not all clients can run all scheduled operations, even though you can define the schedule on the server and associate it with the client. For example, a Macintosh client cannot run a schedule when the action is to restore or retrieve files, or run an executable script. An executable script is also known as a command file, a batch file, or a script on different client operating systems.

IBM Spectrum Protect cannot run multiple schedules concurrently for the same client node.

Privilege class

To define a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.

Syntax

```
Classic client schedule
>>-DEFine SCHedule--domain_name--schedule_name----->
>--+-----+-----+-----+-----+-----+----->
  '-Type----Client-'  '-DESCription--==description-'
```



```

                                +-October---+
                                +-November--+
                                '-December--'

.-DAYOFMonth-----ANY----- .-WEEKofmonth----ANY----- .
>-----+-----+-----+-----+-----+-----+----->
'-DAYOFMonth-----+ANY-+-' '-WEEKofmonth----+ANY-+-'
      '-Day-'                               +-First--+
                                           +-Second-+
                                           +-Third--+
                                           +-FOurth-+
                                           '-Last---'

.-DAYofweek-----ANY----- .
>-----+-----+-----+-----+-----+----->
'-DAYofweek-----+ANY-+-'
      +-WEEKDay---+
      +-WEEKEnd---+
      +-SUnDay----+
      +-Monday----+
      +-TUesday---+
      +-WednesDay--+
      +-THurSday--+
      +-FRiday----+
      '-SATurday--'

.-EXPIration-----Never----- .
>-----+-----+-----+-----+-----+-----><
'-EXPIration-----+Never-+-'
      '-date--'

```

Notes:

1. The OBJECTS parameter is optional when ACTION=INCREMENTAL, but is required for other actions.

Parameters

domain_name (Required)

Specifies the name of the policy domain to which this schedule belongs.

schedule_name (Required)

Specifies the name of the schedule to be defined. You can specify up to 30 characters for the name.

Type=Client

Specifies that a schedule for a client is defined. This parameter is optional.

DESCRiption

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description.

Enclose the description in quotation marks if it contains any blank characters.

ACTION

Specifies the action that occurs when this schedule is processed. Possible values are:

Incremental

Specifies that the schedule backs up all files that are new or that have changed since the last incremental backup.

Incremental also backs up any file for which all existing backups might have expired.

Selective

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

Archive

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

Backup

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

REStore

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

RETRieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

Remember: A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.

IMAGERESTore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.

Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

SUBACTion

You can specify one of the following values:

""

When a null string (two double quotes) is specified with ACTION=BACKUP the backup is an incremental.

FASTBACK

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

SYSTEMSTATE

Specifies that a client Systemstate backup is scheduled.

VApp

Specifies that a client vApp backup is scheduled. A vApp is a collection of pre-deployed virtual machines.

VM

Specifies that a client VMware backup operation is scheduled.

Deploy

Specifies whether to update client workstations with deployment packages that are specified with the OBJECTS parameter. The OBJECTS parameter must contain two specifications, the package files to retrieve and the location from which to retrieve them. Ensure that the objects are in the order *files location*. For example:

```
define schedule standard deploy_1 action=DEPLOY objects=  
"\\IBM_ANR_WIN\c$\tsm\maintenance\client\v6r2\Windows\X32\v620\v6200\  
..\IBM_ANR_WIN"
```

Values for the following options are restricted when you specify ACTION=DEPLOY:

PERUNITS

Specify PERUNITS=ONETIME. If you specify PERUNITS=PERIOD, the parameter is ignored.

DURUNITS

Specify MINUTES, HOURS, or DAYS for the DURUNITS parameter. Do not specify INDEFINITE.

SCHEDSTYLE

Specify the default style, CLASSIC.

The SCHEDULE command fails if the parameters do not conform to the required parameter values, such as the V.R.M.F.

OPTions

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no effect when the client processes the scheduled command:

- MAXCMDRETRIES
- OPTFILE
- QUERYSCHEDPERIOD
- RETRYPERIOD
- SCHEDLOGNAME
- SCHEDMODE
- SERVERNAME

- TCPCLIENTADDRESS
- TCPCLIENTPORT

Windows When you define a scheduler service by using the DSMCUTIL command or the backup-archive client GUI wizard, you specify an options file. You cannot override the options in that options file by issuing the scheduled command. You must modify the options in your scheduler service.

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify `subdir=yes` and `domain all-local -systemobject`, enter:
 - `options='-subdir=yes -domain="all-local -c: -systemobject"'`
- To specify `domain all-local -c: -d:`, enter:
 - `options='-domain="all-local -c: -d:"'`

Windows Tip:

For Windows clients running in batch mode, if the use of quotation marks is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

OBJECTS

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when ACTION=INCREMENTAL. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify ACTION=INCREMENTAL without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:

```
ANS1082E Invalid destination file specification <filespec> entered.
```

- If you specify more than two file specifications, you receive this error:

```
ANS1102E Excessive number of command line arguments passed to the program!
```

When you specify ACTION=ARCHIVE, INCREMENTAL, or SELECTIVE for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

Windows If you are using characters that have a special meaning for Windows users, such as commas, surround the entire argument in two pairs of double quotes, then surround the entire string with single quotes. The following examples show you how to specify some file names:

- To specify `C:\FILE 2`, `D:\GIF FILES`, and `E:\MY TEST FILE`, enter:
 - `OBJECTS='"C:\FILE 2" "D:\GIF FILES" "E:\MY TEST FILE"'`
- To specify `D:\TEST FILE`, enter:
 - `OBJECTS='"D:\TEST FILE"'`
- To specify `D:TEST,FILE`:
 - `OBJECTS='""D:\TEST,FILE""'`

The following examples show how to specify some file names:

- To specify /home/file 2, /home/gif files, and /home/my test file, enter:
 - OBJECTS="/home/file 2" "/home/gif files" "/home/my test file"
- To specify /home/test file, enter:
 - OBJECTS="/home/test file"

Windows

Tip:

For Windows clients running in batch mode, if the use of double quotes is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when IBM Spectrum Protect processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the STARTTIME parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00. If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00.
NOW-HH:MM or - HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00.

DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

Tip: Define schedules with durations longer than 10 minutes. Doing this will give the IBM Spectrum Protect scheduler enough time to process the schedule and prompt the client.

DURUnits

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is HOURS.

Use this parameter with the DURATION parameter to specify how long the startup window remains open to process the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. You can specify one of the following values:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE, unless you specify PERUNITS=ONETIME. The INDEFINITE value is not allowed with enhanced schedules.

MAXRUNtime

Specifies the maximum run time, which is the number of minutes during which all client sessions that are started by the scheduled operation should be completed. If sessions are still running after the maximum run time, the server issues a warning message, but the sessions continue to run.

Tip: The maximum run time is calculated from the beginning of the startup window and not from the time that sessions start within the startup window.

Restrictions:

- The value of the parameter is not distributed to servers that are managed by an enterprise configuration manager.
- The value of the parameter is not exported by the EXPORT command.

The parameter is optional. You can specify a number in the range 0-1440. The default value is 0. A value of 0 means that the maximum run time is indefinite, and no warning message is issued. The maximum run time must be greater than the startup window duration, which is defined by the DURATION and DURUNITS parameters.

For example, if the start time of a scheduled operation is 9:00 PM, and the duration of the startup window is 2 hours, the startup window is 9:00 PM - 11:00 PM. If the maximum run time is 240 minutes, that is, 4 hours, all client sessions for this operation should be completed by 1:00 AM. If one or more sessions are still running after 1:00 AM, the server issues a warning message.

Tip: Alternatively, you can specify a *Run time alert* value of 1:00 AM in the IBM Spectrum Protect Operations Center.

SCHEDStyle

This parameter is optional. SCHEDSTYLE defines either the interval between times when a schedule can run, or the days on which it runs. The default is the classic syntax.

Possible values are:

Classic

The parameters for the Classic syntax are: PERIOD, PERUNITS, and DAYOFWEEK. You cannot use these parameters: MONTH, DAYOFMONTH, and WEEKOFMONTH.

Enhanced

The parameters for the Enhanced syntax are: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. You cannot use these parameters: PERIOD and PERUNITS.

PERIOD

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is DAYS.

Use this parameter with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. You can specify one of the following values:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify PERUNITS=MONTHS, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

When you specify PERUNITS=YEARS, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the PERIOD parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the DAYofweek parameter, depending on whether the schedule style was defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. DAYOFWEEK must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

Possible values for the DAYofweek parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

Sunday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

Tuesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAaturday

Specifies that the startup window begins on Saturday.

MONTH

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY, which means that the schedule runs during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, and so on. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs on each of the specified days of the month. If multiple values resolve to the same day, the schedule runs only once that day.

The default value is ANY. ANY means that the schedule runs on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule runs only once during that week.

The default value is ANY. ANY means that the schedule runs during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXPIRATION

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. You can specify one of the following values:

Never

Specifies that the schedule never expires.

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Define a schedule for a monthly incremental backup

Define a schedule named MONTHLY_BACKUP that initiates an incremental backup of all associated nodes. Specify the start date as Tuesday, May 1, 2001. This date does not match the specified day of the week (Sunday), so the initial startup window begins on the first Sunday after May 1, 2001 (05/01/2001). The startup windows for this schedule extend from 01:00 through 03:00. This monthly schedule initiates backup of c: and d: file spaces for all associated nodes.

```
define schedule standard monthly_backup
description="Monthly Backup of c: and d: drives"
objects="c:\* d:\*"
startdate=05/01/2001 starttime=01:00
duration=2 durunits=hours period=1
perunits=months dayofweek=sunday
```

Example: Define a schedule for a weekly incremental backup

Define a schedule named WEEKLY_BACKUP that initiates an incremental backup of all associated nodes. The initial startup window for this schedule extends from 23:00 on Saturday, June 7, 1997 (06/07/1997), to 03:00 on Sunday, June 8, 1997 (06/08/1997). Subsequent windows begin at 23:00, every Saturday. No messages are returned to the client node when this schedule is run.

```
define schedule employee_records weekly_backup
startdate=06/07/1997 starttime=23:00 duration=4
durunits=hours perunits=weeks
dayofweek=saturday options=-quiet
```

Example: Define a schedule that archives a specific directory every quarter

Define a schedule that archives specific files quarterly on the last Friday of the month.

```
define schedule employee_records quarterly_archive
starttime=20:00 action=archive
object=/home/employee/records/*
duration=1 durunits=hour schedstyle=enhanced
month=mar,jun,sep,dec weekofmonth=last dayofweek=fri
```

DEFINE SCHEDULE (Define a schedule for an administrative command)

Use the DEFINE SCHEDULE command to create a new schedule for processing an administrative command.

You can include scripts in an administrative command schedule so the commands are processed automatically.

Note:

1. You cannot schedule the MACRO command or the QUERY ACTLOG command.

- If you are scheduling a command that specifies the WAIT parameter, the parameter must be set to YES in order for the process to provide a return code to the session that started it. For more information about the WAIT parameter, see Server command processing.

Privilege class

To define an administrative command schedule, you must have system privilege.

Syntax

```
Classic administrative schedule
>>-DEFine SChedule--schedule_name----->
>--+-----+--CMD---command----->
  '-Type---Administrative-'
  .-ACTIVE---No-.
>--+-----+-----+----->
  '-ACTIVE---Yes-' '-DESCRiption---description-'
  .-PRIority---5----- .-STARTDate---current_date-.
>--+-----+-----+----->
  '-PRIority---number-' '-STARTDate---date-----'
  .-STARTTime---current_time-. .-DURation---1-----
>--+-----+-----+----->
  '-STARTTime---time-----' '-DURation---number-'
  .-DURUnits---Hours----- .-MAXRUNtime---0-----
>--+-----+-----+----->
  '-DURUnits---+Minutes-----' '-MAXRUNtime---number-'
                    +-Hours-----+
                    +-Days-----+
                    '-INDefinite-'
  .-SCHEDStyle---Classic-. .-PERiod---1-----
>--+-----+-----+----->
  '-SCHEDStyle---Classic-' '-PERiod---number-'
  .-PERUnits---Days-----
>--+-----+-----+----->
  '-PERUnits---+Hours---+'
                    +-Days---+
                    +-Weeks---+
                    +-Months---+
                    +-Years---+
                    '-Onetime-'
  .-DAYofweek---ANY-----
>--+-----+-----+----->
  '-DAYofweek---+ANY-----+'
                    +-WEEKDay---+
                    +-WEEKEnd---+
                    +-SUNDay---+
                    +-Monday---+
                    +-TUESday---+
                    +-WEDnesday---+
                    +-THURsday---+
                    +-FRIDay---+
                    '-SATurday--'
  .-EXPIration---Never-----
>--+-----+-----+----->>
  '-EXPIration---+Never---+'
                    '-date--'
```

Syntax

Specifies the name of the schedule to be defined. You can specify up to 30 characters for the name.

Type=Administrative

Specifies that a schedule for an administrative command is defined. This parameter is optional. An administrative command is assumed if the CMD parameter is specified.

CMD (Required)

Specifies the administrative command to schedule for processing. The maximum length of the command is 512 characters. Enclose the administrative command in quotation marks if it contains any blank characters.

Restriction: You cannot specify redirection characters with this parameter.

ACTIVE

Specifies whether IBM Spectrum Protect processes an administrative command schedule when the startup window occurs. This parameter is optional. The default is NO. The administrative command schedule must be set to the active state with the UPDATE SCHEDULE command so that IBM Spectrum Protect can process the schedule. Possible values are:

YES

Specifies that IBM Spectrum Protect processes an administrative command schedule when the startup window begins.

NO

Specifies that IBM Spectrum Protect does not process an administrative command schedule when the startup window begins.

DEScRiption

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains any blank characters.

PRIOrity

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when IBM Spectrum Protect processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the STARTTIME parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00. If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00.

DURATION

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

DURUNITS

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is HOURS.

Use this parameter with the DURATION parameter to specify how long the startup window remains open to process the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. You can specify one of the following values:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE, unless you specify PERUNITS=ONETIME. The INDEFINITE value is not allowed with enhanced schedules.

MAXRUNtime

Specifies the maximum run time, which is the number of minutes during which server processes that are started by the scheduled commands must be completed. If processes are still running after the maximum run time, the central scheduler cancels the processes.

Tips:

- The processes might not end immediately when the central scheduler cancels them; they end when they register the cancellation notification from the central scheduler.
- The maximum run time is calculated beginning from when the server process starts. If the schedule command starts more than one process, each process maximum run time is calculated from when the process starts.
- This parameter does not apply to some processes, such as duplicate-identification processes, which can continue to run after the maximum run time.
- This parameter does not apply if the scheduled command does not start a server process.
- Another cancel time might be associated with some commands. For example, the MIGRATE STGPOOL command can include a parameter that specifies the length of time that the storage pool migration runs before the migration is

automatically canceled. If you schedule a command for which a cancel time is defined, and you also define a maximum run time for the schedule, the processes are canceled at whichever cancel time is reached first.

Restrictions:

- The value of the parameter is not distributed to servers that are managed by an enterprise configuration manager.
- The value of the parameter is not exported by the EXPORT command.

The parameter is optional. You can specify a number in the range 0-1440. The default value is 0. A value of 0 means that the maximum run time is indefinite, and the central scheduler does not cancel processes. The maximum run time must be greater than the startup window duration, which is defined by the DURATION and DURUNITS parameters.

For example, if the start time of a scheduled command is 9:00 PM, and the duration of the startup window is 2 hours, the startup window is 9:00 PM - 11:00 PM. If the maximum run time is 240 minutes, that is, 4 hours, all applicable server processes that are started by the command must be completed by 1:00 AM. If one or more applicable processes are still running after 1:00 AM, the central scheduler cancels the processes.

Tip: Alternatively, you can specify an *end time* of 1:00 AM in the IBM Spectrum Protect Operations Center.

SCHEDStyle

This parameter is optional. SCHEDSTYLE defines either the interval between times when a schedule should run, or the days on which it should run. The style can be either classic or enhanced. The default is the classic syntax.

For classic schedules, these parameters are allowed: PERIOD, PERUNITS, and DAYOFWEEK. Not allowed for classic schedules are: MONTH, DAYOFMONTH, and WEEKOFMONTH.

For enhanced schedules, these parameters are allowed: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. These parameters are not allowed: PERIOD and PERUNITS.

PERiod

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is DAYS.

Use this parameter with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. You can specify one of the following values:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify PERUNITS=MONTHS, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

When you specify PERUNITS=YEARS, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the PERIOD parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the DAYofweek parameter, depending on whether the schedule style was defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. DAYOFWEEK must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

Possible values for the DAYofweek parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUnDay

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

TUesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

THursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAturday

Specifies that the startup window begins on Saturday.

MONth

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY. This means the schedule will run during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, etc. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will

run on each of the specified days of the month. If multiple values resolve to the same day, the schedule will run only once that day.

The default value is ANY. This means the schedule will run on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule will run only once during that week.

The default value is ANY, meaning the schedule will run during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXpiration

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. You can specify one of the following values:

Never

Specifies that the schedule never expires.

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Define a schedule to back up the primary storage pool every two days

Define a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. The backup runs at 8 p.m. every two days.

```
define schedule backup_archivepool type=administrative
cmd="backup stgpool archivepool recoverypool"
active=yes starttime=20:00 period=2
```

Example: Define a schedule to back up the primary storage pool twice a month

Define a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. Select an enhanced schedule and run on the first and fifteenth day of the month.

```
define schedule backup_archivepool type=administrative
cmd="backup stgpool archivepool recoverypool"
schedstyle=enhanced dayofmonth=1,15
```

DEFINE SCRATCHPADENTRY (Define a scratch pad entry)

Use this command to enter data on a new line in the scratch pad. The scratch pad is a database table that the server hosts. You can use the scratch pad to store diverse information in table format.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine SCRATCHPadentry--major_category--minor_category----->
>--subject--Line-----number--Data---data-----><
```

Parameters

major_category (Required)

Specifies the major category in which data is to be stored. Enter a text string of up to 100 alphanumeric characters. This parameter is case sensitive.

minor_category (Required)

Specifies the minor category in which data is to be stored. Minor categories are sections within major categories. Enter a text string of up to 100 alphanumeric characters. This parameter is case sensitive.

subject (Required)

Specifies the subject under which data is to be stored. Subjects are sections within minor categories. Enter a text string of up to 100 alphanumeric characters. This parameter is case sensitive.

Line (Required)

Specifies the number of the line on which data is to be stored. Lines are sections within subjects. Specify an integer in the range 1 - 1000.

Data (Required)

Specifies the data to be stored on the line. You can enter up to 1000 characters. Enclose the data in quotation marks if the data contains one or more blanks. The data is case sensitive.

Example: Define a scratch pad entry

Enter the vacation dates of an administrator, Jane, in a table that stores information about the location of all administrators.

```
define scratchpadentry admin_info location jane line=2 data=
"Out of the office from 1-15 Nov."
```

Related commands

Table 1. Commands related to DEFINE SCRATCHPADENTRY

Command	Description
DELETE SCRATCHPADENTRY	Deletes a line of data from the scratch pad.
QUERY SCRATCHPADENTRY	Displays information that is contained in the scratch pad.
SET SCRATCHPADRETENTION	Specifies the amount of time for which scratch pad entries are retained.
UPDATE SCRATCHPADENTRY	Updates data on a line in the scratch pad.

DEFINE SCRIPT (Define an IBM Spectrum Protect script)

Use this command to define an IBM Spectrum Protect™ script or to create a new IBM Spectrum Protect script by using the contents from another script.

The first line for the script can be defined with this command. To add subsequent lines to the script, use the UPDATE SCRIPT command.

Tips:

- When routing commands inside scripts, enclose the server or server group in parentheses and omit the colon. Otherwise, if the syntax includes a colon, the command is not routed when the RUN command is issued. Instead, the command runs only on the server from which the RUN command is issued.
- You cannot redirect the output of a command within an IBM Spectrum Protect script. Instead, run the script and then specify command redirection. For example, to direct the output of script1 to the c:\temp\test.out directory, run the script and specify command redirection as in the following example:

```
run script1 > c:\temp\test.out
```

Privilege class

To issue this command, you must have operator, policy, storage, or system privilege.

Syntax

```
>>--DEFine SCRipt--script_name----->
```

```

                .-Line-----001----.
>--+--command_line--+-----+----->
|                '-Line ----number-' |
| '-File-----file_name-----' |
>--+-----+----->>
| '-DESCRiption-----description-'

```

Parameters

script_name (Required)

Specifies the name of the script to be defined. You can specify up to 30 characters for the name.

command_line

Specifies the first command to be processed in a script. You must specify either this parameter (and optionally, the LINE parameter) or the FILE parameter.

The command that you specify can include substitution variables and can be continued across multiple lines if you specify a continuation character (-) as the last character in the command. Substitution variables are specified with a '\$' character, followed by a number that indicates the value of the parameter when the script is processed. You can specify up to 1200 characters for the command line. Enclose the command in quotation marks if it contains blanks.

You can run commands serially, in parallel, or serially and in parallel by specifying the SERIAL or PARALLEL script commands for the COMMAND_LINE parameter. You can run multiple commands in parallel and wait for them to complete before you proceed to the next command. Commands run serially until the parallel command is encountered.

Conditional logic flow statements can be used. These statements include IF, EXIT, and GOTO.

Line

Specifies the line number for the command line. Because commands are specified in multiple lines, line numbers are used to determine the order for processing when the script is run. The first line, or line 001 is the default. This parameter is optional.

File

Specifies the name of the file whose contents are read into the script to be defined. The file must reside on the server where this command is running. If you specify the FILE parameter, you cannot specify a command line or line number.

You can create a script by querying another script and specifying the FORMAT=RAW and OUTPUTFILE parameters. The output from querying the script is directed to a file you specify with the OUTPUTFILE parameter. To create the new script, the contents of the script to be defined are read in from the file you specified with the OUTPUTFILE parameter.

DESCRiption

Specifies a description for the script. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters. This parameter is optional.

Example: Write a script to display AIX clients

Define a script that displays all AIX® clients.

```

define script qaixc "select node_name from nodes where platform_name='AIX'"
desc='Display aix clients'

```

Example: Write and run a script to route a command to a server group

Define and run a script that routes the QUERY STGPOOL command to a server group named DEV_GROUP.

```

define script qu_stg "(dev_group) query stgpool"

run qu_stg

```

Example: Create a script from an existing script

Define a script whose command lines are read in from a file that is named MY.SCRIPT and name the new script AGADM. The file must be on the server, and be read by the server.

```
define script agadm file=my.script
```

Related commands

Table 1. Commands related to DEFINE SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

Related concepts:

Using logic flow statements in a script

Related tasks:

Defining a server script

Running commands in parallel or serially

Performing tasks concurrently on multiple servers


Related reference:

Return codes for use in IBM Spectrum Protect scripts

DEFINE SERVER (Define a server for server-to-server communications)

Use this command to define a server to use functions such as virtual volumes, node replication, command routing, and LAN-free data movement, among others.

Use this command to define a server for the following functions:

- Enterprise configuration
- Enterprise event logging
- Command routing
- Virtual volumes
- LAN-free data movement
- Node replication
-  Data movement by using z/OS® media server
- Status monitoring of remote servers
- Alert monitoring of remote servers
- Server-to-server export

If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP-authenticated passwords. Data that is replicated from a node that authenticates with an LDAP directory server is inaccessible if the target replication server is not properly configured. If your target replication server is not configured, replicated data from an LDAP node can make it to the target server. But the target replication server must be configured to use LDAP if you want to access the data.

The use of virtual volumes is not supported when the source server and the target server are on the same IBM Spectrum Protect™ server.

This command also is used to define an IBM Spectrum Protect storage agent as if it were a server.

Privilege class

To issue this command, you must have system privilege.

Syntax

For:

- Command routing
- Status monitoring of remote servers


- Alert monitoring of remote servers
- Server-to-server export

Tip: Command routing uses the ID and the password of the administrator who is issuing the command.

```
>>-DEfINE--SERver--server_name--HLAddress-----ip_address----->
>--LLAddress-----tcp_port--+-----+----->
      '-COMMmethod-----TCPIP-'
>--+-----+-----+-----+----->
      '-URL-----url-'   '-DEScRiption-----description-'
      .-SSL-----No-----.
>--+-----+-----+-----+----->
      '-SSL-----+No--+-'
          '-Yes-'
      .-SESSiONSECurity-----TRANSitional-----.
>--+-----+-----+-----+----->>
      '-SESSiONSECurity-----+STRict-----+-'
          '-TRANSitional-'
```

Syntax

For:

- Enterprise configuration
- Enterprise event logging
- Storage agent
- Node replication source and target servers
-  z/OS media server

```
>>-DEfINE--SERver--server_name--SERVERPAssword-----password----->
>--HLAddress-----ip_address--LLAddress-----tcp_port----->
>--+-----+-----+-----+----->
      '-COMMmethod-----TCPIP-'   '-URL-----url-'
>--+-----+-----+-----+----->
      '-DEScRiption-----description-'
      (1)
      .-CROSSDEfINE-----No----- (2)
>--+-----+-----+-----+----->
      '-CROSSDEfINE-----+No--+-'
          '-Yes-'
      .-VALIdateprotocol-----No-----.   .-SSL-----No-----.
>--+-----+-----+-----+----->
      '-VALIdateprotocol-----+No--+-'   '-SSL-----+No--+-'
          '-All-'   '-Yes-'
      .-SESSiONSECurity-----TRANSitional-----.
>--+-----+-----+-----+----->
      '-SESSiONSECurity-----+STRict-----+-'
          '-TRANSitional-'
      .-TRANSFERMethod-----TcpiP-----.
>--+-----+-----+-----+----->>
      '-TRANSFERMethod-----+TcpiP-----+-'
          | (3) |
          '-Fasp-----'
```

Notes:

1. The CROSSDEFINE parameter does not apply to storage agent definitions.

2. The VALIDATEPROTOCOL parameter is deprecated and applies only to storage agent definitions.
3. **Linux** The TRANSFERMETHOD parameter is available only on Linux x86_64 operating systems.

Syntax for virtual volumes

```
>>-DEFine--SERver--server_name--PAssword---password----->
>--HLAddress---ip_address--LLAddress---tcp_port----->
>--+-----+-----+-----+-----+----->
  '-COMMmethod---TCPIP-' '-URL---url-'
>--+-----+-----+-----+-----+----->
  '-DELgraceperiod---days-' '-NODEName---node_name-'
                                     .-SSL---No-----
>--+-----+-----+-----+-----+----->
  '-DESCRiption---description-' '-SSL---+No---+'
                                     '-Yes-'
                                     .-SESSIONSECurity---TRANSitional-----
>--+-----+-----+-----+-----+-----><
  '-SESSIONSECurity---+STRict-----+'
                                     '-TRANSitional-'
```

Parameters

server_name (Required)

Specifies the name of the server. This name must be unique on the server. The maximum length of this name is 64 characters.

For server-to-server event logging, library sharing, and node replication, you must specify a server name that matches the name that was set by issuing the SET SERVERNAME command at the target server.

PAssword

Specifies the password that is used to sign on to the target server for virtual volumes. If you specify the NODENAME parameter, you must specify the PASSWORD parameter. If you specify the PASSWORD parameter but not the NODENAME parameter, the node name defaults to the server name that is specified with the SET SERVERNAME command. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

SERVERPAssword

Specifies the password of the server that you are defining. This password must match the password that is set by the SET SERVERPASSWORD command. This parameter is required for enterprise configuration and server-to-server event logging functions. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

HLAddress (Required)

Specifies the IP address (in dotted decimal format) of the server.

Do not use the loopback address as the value of this parameter. Virtual volumes are not supported when the source server and the target server are the same IBM Spectrum Protect server.

LLAddress (Required)

Specifies the low-level address of the server. This address is usually the same as the address in the TCPPOrt server option of the target server. When SSL=YES, the port must already be designated for SSL communications on the target server.

COMMmethod

Specifies the communication method that is used to connect to the server. This parameter is optional.

URL

Specifies the URL address of this server. The parameter is optional.

DELgraceperiod

Specifies a number of days that an object remains on the target server after it was marked for deletion. You can specify a value 0 - 9999. The default is 5. This parameter is optional.

NODEName

Specifies a node name to be used by the server to connect to the target server. This parameter is optional. If you specify the NODENAME parameter, you must also specify the PASSWORD parameter. If you specify the PASSWORD parameter but not

the NODENAME parameter, the node name defaults to the server name specified with the SET SERVERNAME command.

DEscription

Specifies a description of the server. The parameter is optional. The description can be up to 255 characters. Enclose the description in quotation marks if it contains blank characters.

CROSSDEFine

Specifies whether the server that is running this command defines itself to the server that is being specified by this command. This parameter is optional.

AIX | **Linux** | **Windows** Important: This parameter does not apply to storage agent definitions. If this parameter is included, you must also issue the SET SERVERNAME, SET SERVERPASSWORD, SET SERVERHLADDRESS, SET CROSSDEFINE, and SET SERVERLLADDRESS commands. The default is NO.

Remember:

- For replication operations, the names of the source and target replication servers must match the names that you specify in this command.
- CROSSDEFINE can be used with SSL=YES if all of the conditions that are specified for the SSL=YES parameter are in place on the source and target server.

You can specify one of the following values:

No

Cross definition is not completed.

Yes

Cross definition is completed.

VALIDateprotocol (deprecated)

Specifies whether a cyclic redundancy check validates the data that is sent between the storage agent and IBM Spectrum Protect server. The parameter is optional. The default is NO.

Important: Beginning with IBM Spectrum Protect Version 8.1.2 and Tivoli® Storage Manager Version 7.1.8, validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The VALIDATEPROTOCOL parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

SSL

Specifies the communication mode of the server. The default is NO.

Important: Beginning in IBM Spectrum Protect V8.1.2 and Tivoli Storage Manager V7.1.8, the SSL parameter uses SSL to encrypt some communication with the specified server even if SSL=NO.

The following conditions and considerations apply when you specify the SSL parameter:

- Before you start the servers, self-signed certificates of the partner servers must be in the key database file (cert.kdb) of each of the servers.
- You can define multiple server names with different parameters for the same target server.
- Storage agents can issue the DSMSTA SETSTORAGESERVER command and include the SSL parameter to create the key database.

You can specify one of the following values:

No

Specifies an SSL session for all communication with the specified server, except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure.

Yes

Specifies an SSL session for all communication with the specified server, even when the server is sending and receiving object data.

SESSIONSECurity

Specifies whether the server that you are defining must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRict

Specifies that the strictest security settings are enforced for the server that you are defining. The STRICT value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the specified server and an IBM Spectrum Protect server.

To use the STRICT value, the following requirements must be met to ensure that the specified server can authenticate with the IBM Spectrum Protect server:

- Both the server that you are defining and the IBM Spectrum Protect server must be using IBM Spectrum Protect software that supports the SESSIONSECURITY parameter.
- The server that you are defining must be configured to use the TLS 1.2 protocol for SSL sessions between itself and the IBM Spectrum Protect server.

Servers set to STRICT that do not meet these requirements are unable to authenticate with the IBM Spectrum Protect server.

TRANSitional

Specifies that the existing security settings are enforced for the server. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the STRICT value.

If SESSIONSECURITY=TRANSITIONAL and the server has never met the requirements for the STRICT value, the server will continue to authenticate by using the TRANSITIONAL value. However, after a server meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. Then, the server can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT. In addition, after a server successfully authenticates by using a more secure communication protocol, the server can no longer authenticate by using a less secure protocol. For example, if a server that is not using SSL is updated and successfully authenticates by using TLS 1.2, the server can no longer authenticate by using no SSL protocol or TLS 1.1. This restriction also applies when you use functions such as virtual volumes, command routing, or server-to-server export, when a node or administrator authenticates to the IBM Spectrum Protect server as a node or administrator from another server.

Linux TRANSFERMethod

Specifies the method that is used for server-to-server data transfer. This parameter is optional. You can specify one of the following values:

Tcpip

Specifies that TCP/IP is used to transfer data. This is the default.

Fasp

Specifies that Aspera® Fast Adaptive Secure Protocol (FASP®) technology is used to transfer data. Aspera FASP technology can help you optimize data transfer in a wide area network (WAN).

Restrictions:

- Before you enable Aspera FASP technology, determine whether the technology is appropriate for your system environment and install the appropriate licenses. For instructions, see Determining whether Aspera FASP technology can optimize data transfer in your system environment. If the licenses are missing or expired, data transfer operations fail.
- If WAN performance meets your business needs, do not enable Aspera FASP technology.
- If you specify TRANSFERMETHOD=FASP on the PROTECT STGPOOL or REPLICATE NODE command, that value overrides the TRANSFERMETHOD parameter on the DEFINE SERVER and UPDATE SERVER commands.

Example: Set up two servers to use SSL to communicate (manual configuration)

Tip: If both servers are using IBM Spectrum Protect V8.1.2 or later software or Tivoli Storage Manager V7.1.8 software, SSL is automatically configured between the servers and manual configuration is not required.

If both servers are not using V7.1.8 or V8.1.2 or later software, you must manually configure the two servers to use SSL to communicate.

The server addresses are as follows:

- ServerA is at `bfa.tucson.ibm.com`
- ServerB is at `bfb.tucson.ibm.com`

Complete the following steps to set up the two servers for SSL:

1. Specify option TCPPORT 1500 for both servers in the dsmserv.opt option file.
2. Start both servers.
3. Shut down both servers to import the cert256 partner certificate. For ServerA, the certificate is in the /tsma instance directory. For ServerB, the certificate is in the /tsmb instance directory.
4. Start both servers. The /tsma/cert256.arm file is copied to /tsmb/cert256.bfa.arm on the bfb.tucson.ibm.com address. The /tsmb/cert256.arm file is copied to /tsmb/cert256.bfb.arm on the bfa.tucson.ibm.com address.

5. Issue the following command:

- o From ServerA:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label "bfb" -file /tsma/cert256.bfb.arm
```

- o From ServerB:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii  
-label "bfa" -file /tsmb/cert256.bfa.arm
```

From each server, you can view the certificates in the key database by issuing the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

6. Restart the servers.

7. Issue the appropriate DEFINE SERVER command. For ServerA, issue the following example command:

```
DEFINE SERVER BFB hla=bfb.tucson.ibm.com lla=1542  
serverpa=passwordforbfb SSL=YES
```

For ServerB, issue the following example command:

```
DEFINE SERVER BFA hla=bfa.tucson.ibm.com lla=1542  
serverpa=passwordforbfa SSL=YES
```

If you do not use SSL, issue the following example DEFINE SERVER command on ServerA:

```
DEFINE SERVER BFBTCP hla=bfb.tucson.ibm.com lla=1500  
serverpa=passwordforbfb SSL=NO
```

If you do not use SSL, issue the following example DEFINE SERVER command on ServerB:

```
DEFINE SERVER BFATCP hla=bfa.tucson.ibm.com lla=1500  
serverpa=passwordforbfa SSL=NO
```

Example: Define a server to communicate with another server by using strict session security

Define a server name of SERVER1 to use the strictest security settings to authenticate with the IBM Spectrum Protect server.

```
define server server1 sessionsecurity=strict
```

Example: Define a target server

A target server has a high-level address of 9.116.2.67 and a low-level address of 1570. Define that target server to the source server, name the target server SERVER2, and set the password to SECRETPASSWORD. Specify that objects remain on the target server for seven days after they are marked for deletion.

```
define server server2 password=secretpassword  
hladdress=9.116.2.67 lladdress=1570 delgraceperiod=7
```

Example: Define a server to receive commands from other servers

Define a server that can receive commands that are routed from other servers. Name the server WEST_COMPLEX. Set the high-level address to 9.172.12.35, the low-level address to 1500, and the URL address to http://west_complex:1580/.

```
define server west_complex  
hladdress=9.172.12.35 lladdress=1500  
url=http://west_complex:1580/
```

Example: Cross-define two servers

Use cross definition to define SERVER_A and SERVER_B.

1. On SERVER_B, specify the server name, password, and high- and low-level addresses of SERVER_B. Specify that cross defining is allowed.

```
set servername server_b  
set serverpassword mylifepwd  
set serverhladdress 9.115.20.80
```



```
set serverlladdress 1860
set crossdefine on
```

2. On SERVER_A, specify the server name, password, and high- and low-level addresses of SERVER_A.









```
set servername server_a
set serverpassword yourlifepwd
set serverhladdress 9.115.20.97
set serverlladdress 1500
```

3. On SERVER_A, define SERVER_B:

```
define server server_b hladdress=9.115.20.80 lladdress=1860
serverpassword=mylifepwd crossdefine=yes
```

Related commands

Table 1. Commands related to DEFINE SERVER

Command	Description
DEFINE DEVCLASS	Defines a device class.
  DEFINE PATH	  Define a path when the destination is a z/OS media server.
DELETE DEVCLASS	Deletes a device class.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
DELETE SERVER	Deletes the definition of a server.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY SERVER	Displays information about servers.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERPASSWORD	Specifies the server password.
SET REPLSERVER	Specifies a target replication server.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE NODE	Changes the attributes that are associated with a client node.
  UPDATE PATH	  Define a path when the destination is a z/OS media server.
UPDATE SERVER	Updates information about a server.

DEFINE SERVERGROUP (Define a server group)

Use this command to define a server group. With a server group, you can route commands to multiple servers by specifying only the group name. After you define the server group, add servers to the group by using the DEFINE GRPMEMBER command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine SERVERGroup--group_name----->
>--+-----+-----><
  '-DESCription----description-'
```

Parameters

group_name (Required)

Specifies the name of the server group. The maximum length of the name is 64 characters.

DESCRIPTION

Specifies a description of the server group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a server group

Define a server group named WEST_COMPLEX.

```
define servergroup west_complex
```

Related commands

Table 1. Commands related to DEFINE SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DEFINE SPACETRIGGER (Define the space trigger)

Use this command to define settings for triggers that determine when and how the server prepares extra space when predetermined thresholds are exceeded in storage pools that use FILE and DISK device classes. Space triggers are not enabled for storage pools with a parameter RECLAMATIONTYPE=SNAPLOCK.

The IBM Spectrum Protect™ server allocates more space when space utilization reaches a specified value. After allocating more space, the server either adds the space to the specified pool (random-access or sequential-access disk).

Important: Space trigger functions and storage pool space calculations take into account the space remaining in each directory. An inaccurate calculation can result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled.

For example, if you specify multiple directories for a device class and the directories reside in the same file system, the server calculates space by adding values representing the space remaining in each directory. These space calculations are inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the directory that is specified for the device class and run out of space prematurely.

To prevent possible problems and ensure an accurate calculation, you associate each directory with a separate file system. If a trigger becomes disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by specifying the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```

      .-Fullpct----80-----.
>>-DEFine SPACETrigger---STG---+-----+----->
      '-Fullpct---percent-'

      .-SPACEexpansion---20-----.
>--+-----+----->
      '-SPACEexpansion---percent-'

>--+-----+----->
      '-EXPansionprefix---prefix-'

>--+-----+-----><
      '-STGPOOL---storage_pool_name-'

```

Parameters

STG

Specifies a storage pool space trigger.

Fullpct

This parameter specifies the utilization percentage of the storage pool. This parameter is optional. Specify an integer value 0 - 99. The default is 80. A value of zero (0) disables the space trigger. When this value is exceeded, the space trigger creates new volumes. Exceeding the threshold might not cause new volumes to be created until the next space request is made.

You can determine storage pool utilization by issuing the `QUERY STGPOOL` command with `FORMAT=DETAILED`. The percentage of storage pool utilization is displayed in the field "Space Trigger Util." The calculation for this percentage does not include potential scratch volumes. The calculation for the percentage utilization that is used for migration and reclamation, however, does include potential scratch volumes.

SPACEexpansion

For sequential-access FILE-type storage pools, this parameter is used in determining the number of additional volumes that are created in the storage pool. This parameter is optional. The default is 20. Volumes are created using the `MAXCAPACITY` value from the storage pool's device class. For random-access DISK storage pools, the space trigger creates a single volume using the `EXPANSIONPREFIX`.

EXPansionprefix

For random-access DISK storage-pools, this parameter specifies the prefix that the server uses to create new storage pool files. This parameter is optional and applies only to random-access DISK device classes. The default prefix is the server installation path.

The prefix can include one or more directory separator characters, for example:

AIX | **Linux**

```
/opt/tivoli/tsm/server/bin/
```

Windows

```
c:\program files\tivoli\tsm\
```

AIX

Linux

You can specify up to 250 characters. If you specify an invalid prefix, automatic expansion can fail.

Windows

You can specify up to 200 characters. If you specify an invalid prefix, automatic expansion can fail. If the server is running as a Windows service, the default prefix is the `c:\wnnt\system32` directory.

This parameter is not valid for space triggers for sequential-access FILE storage pools. Prefixes are obtained from the directories that are specified with the associated device class.

STGPOOL

Specifies the storage pool that is associated with this space trigger. This parameter is optional for storage pool space triggers. If you specify the STG parameter but not the STGPOOL parameter, one space trigger is created that applies to all random-access DISK and sequential-access FILE storage pools that do not have a specific space trigger.

This parameter does not apply to storage pools with the parameter RECLAMATIONTYPE=SNAPLOCK.

Example: Define a space trigger to increase storage pool space 25 percent

Set up a storage pool space trigger for increasing the amount of space in a storage pool by 25 percent when it is filled to 80 percent utilization of existing volumes. Space is created in the directories associated with the device class.

```
define spacetrigger stg spaceexpansion=25 stgpool=file
```

Example: Define a space trigger to increase storage pool space 40 percent

Set up a space trigger for the WINPOOL1 storage pool to increase the amount of space in the storage pool by 40 percent when it is filled to 80 percent utilization of existing volumes.

```
define spacetrigger stg spaceexpansion=40 stgpool=winpool1
```

Related commands

Table 1. Commands related to DEFINE SPACETRIGGER

Command	Description
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE SPACETRIGGER	Deletes the storage pool space trigger.
QUERY SPACETRIGGER	Displays information about a storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)

Use this command to define a new status monitoring threshold.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>>DEFine STATusthreshold--threshold_name--activity----->
    .-Condition----EXists----.
>---+-----+-----+-----+-----+-----+-----+----->
    '-Condition----EXists-+-' '-Value----value-'
        +-GT-----+
        +-GE-----+
        +-LT-----+
        +-LE-----+
        '-Equal--'
```

```

.-Status---Normal-----.
>-----+-----><
'-Status---+Normal---+'
      +-Warning-+
      '-Error---'

```

Parameters

threshold_name (Required)

Specifies the threshold name. The name cannot exceed 48 characters in length.

activity (Required)

Specifies the activity for which you want to create status indicators. Specify one of the following values:

PROCESSSUMMARY

Specifies the number of processes that are currently active.

SESSIONSUMMARY

Specifies the number of sessions that are currently active.

CLIENTSESSIONSUMMARY

Specifies the number of client sessions that are currently active.

SCHEDCLIENTSESSIONSUMMARY

Specifies the number of scheduled client sessions.

DBUTIL

Specifies the database utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

DBFREESPACE

Specifies the free space available in the database in gigabytes.

DBUSEDSPACE

Specifies the amount of database space that is used, in gigabytes.

ARCHIVELOGFREESPACE

Specifies the free space that is available in the archive log, in gigabytes.

STGPOOLUTIL

Specifies the storage pool utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

STGPOOLCAPACITY

Specifies the storage pool capacity in gigabytes.

AVGSTGPOOLUTIL

Specifies the average storage pool utilization percentage across all storage pools. The default warning threshold value is 80%, and the default error threshold value is 90%.

TOTSTGPOOLCAPACITY

Specifies the total storage pool capacity in gigabytes for all available storage pools.

TOTSTGPOOLS

Specifies the number of defined storage pools.

TOTRWSTGPOOLS

Specifies the number of defined storage pools that are readable or writeable.

TOTNOTRWSTGPOOLS

Specifies the number of defined storage pools that are not readable or writeable.

STGPOOLINUSEANDDEFINED

Specifies the total number of defined volumes that are in use.

ACTIVELOGUTIL

Specifies the current percent utilization of the active log. The default warning threshold value is 80%, and the default error threshold value is 90%.

ARCHLOGUTIL

Specifies the current utilization of the archive log. The default warning threshold value is 80%, and the default error threshold value is 90%.

CPYSTGPOOLUTIL

Specifies the percent utilization for a copy storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

PMRYSTGPOOLUTIL

Specifies the percent utilization for a primary storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

DEVCLASSPCTDRVOFFLINE

- Specifies the percent utilization of drives that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- DEVCLASSPCTDRVPOLLING**
Specifies the drives polling, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- DEVCLASSPCTLIBPATHSOFFLINE**
Specifies the library paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- DEVCLASSPCTPATHSOFFLINE**
Specifies the percentage of device class paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- DEVCLASSPCTDISKSNOTRW**
Specifies the percentage of disks that are not writable for the disk device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- DEVCLASSPCTDISKSUNAVAILABLE**
Specifies the percentage of the disk volumes that are unavailable, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.
- FILEDEVCLASSPCTSCRUNALLOCATABLE**
Specifies the percentage of scratch volumes that the server cannot allocate for a given non-shared file device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

Condition

Specifies the condition that is used to compare the activity output to the specified value. The default value is EXISTS. Specify one of the following values:

EXists

Creates a status monitoring indicator if the activity exists.

GT

Creates a status monitoring indicator if the activity outcome is greater than the specified value.

GE

Creates a status monitoring indicator if the activity outcome is greater than or equal to the specified value.

LT

Creates a status monitoring indicator if the activity outcome is less than the specified value.

LE

Creates a status monitoring indicator if the activity outcome is less than or equal to the specified value.

EQual

Creates a status monitoring indicator if the activity outcome is equal to the specified value.

Value (Required)

Specifies the value that is compared with the activity output for the specified condition. You must specify this parameter, unless CONDITION is set to EXISTS. You can specify an integer in the range 0 - 999999999999999.

Status

Specifies that the status indicator created in status monitoring if the condition that is being evaluated passes. This optional parameter has a default value of NORMAL. Specify one of the following values:

Normal

Specifies that the status indicator has a normal status value.

Warning

Specifies that the status indicator has a warning status value.

Error

Specifies that the status indicator has an error status value.

Define status threshold

Define a status threshold for average storage pool utilization percentage by issuing the following command:

```
define statusthreshold avgstgpl "AVGSTGPOOLUTIL" value=85
condition=gt status=warning
```

Related commands

Table 1. Commands related to DEFINE STATUSTHRESHOLD

Command	Description
---------	-------------

Command	Description
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

DEFINE STGPOOL (Define a storage pool)

Use this command to define a primary storage pool, copy storage pool, an active-data pool, a directory container storage pool, a container-copy storage pool, or a container storage pool in a cloud environment.

A primary storage pool provides a destination for backup files, archive files, or files that are migrated from client nodes. A copy storage pool provides a destination for copies of files that are in primary storage pools. An active-data pool provides a destination for active versions of backup data that are in primary storage pools. A container storage pool provides a destination for deduplicated files. A cloud storage pool provides storage in a cloud environment. A container-copy storage pool provides a tape copy of a directory-container storage pool. The maximum number of storage pools that you can define for a server is 999.

All volumes in a storage pool belong to the same device class. Random access storage pools use the DISK device type. After you define a random access storage pool, you must define volumes for the pool to create storage space.

Sequential access storage pools use device classes that you define for tape devices, files on disk (FILE device type), and storage on another server (SERVER device type). To create storage space in a sequential access storage pool, you must allow scratch volumes for the pool when you define or update it, or define volumes for the pool after you define the pool. You can also do both.

Restriction: If a client is using the simultaneous-write function and data deduplication, the data deduplication feature is disabled during backups to a storage pool.

The DEFINE STGPOOL command takes seven forms. The syntax and parameters for each form are defined separately.

Table 1. Commands related to DEFINE STGPOOL

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
COPY ACTIVEDATA	Copies active backup data.
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE DEVCLASS	Defines a device class.

Command	Description
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
DELETE STGPOOL	Deletes a storage pool from server storage.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
MOVE MEDIA	Moves storage pool volumes that are managed by an automated library.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY DEVCLASS	Displays information about device classes.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
QUERY STGPOOL	Displays information about storage pools.
RENAME STGPOOL	Renames a storage pool.
REPAIR STGPOOL	Repairs a directory-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SHRED DATA	Manually starts the process of shredding deleted data.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

- **DEFINE STGPOOL (Define a cloud-container storage pool)**
Use this command to define a container storage pool in a cloud environment. This type of storage pool is used for data deduplication. Cloud-container storage pools are not supported on Linux on System z®.
- **DEFINE STGPOOL (Define a directory-container storage pool)**
Use this command to define a directory-container storage pool that is used for data deduplication.
- **DEFINE STGPOOL (Define a container-copy storage pool)**
Use this command to define a container-copy storage pool to hold a copy of data from a directory-container storage pool.
- **DEFINE STGPOOL (Define a primary storage pool assigned to random access devices)**
Use this command to define a primary storage pool that is assigned to random access devices.
- **DEFINE STGPOOL (Define a primary storage pool assigned to sequential access devices)**
Use this command to define a primary storage pool that is assigned to sequential access devices.
- **DEFINE STGPOOL (Define a copy storage pool assigned to sequential access devices)**
Use this command to define a copy storage pool that is assigned to sequential access devices.
- **DEFINE STGPOOL (Define an active-data pool assigned to sequential-access devices)**
Use this command to define an active-data pool assigned to sequential-access devices.

DEFINE STGPOOL (Define a cloud-container storage pool)

Use this command to define a container storage pool in a cloud environment. This type of storage pool is used for data deduplication. Cloud-container storage pools are not supported on Linux on System z®.

Tip: To optimize backup and archive performance, set up one or more local storage directories to temporarily hold data that IBM Spectrum Protect™ is transferring to the cloud. After you use the DEFINE STGPOOL command to define a cloud-container storage pool, use the DEFINE STGPOOLDIRECTORY command to assign local storage directories to the cloud-container storage pool. For more information, see Optimizing performance for cloud object storage.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--STGType---Cloud----->
    .-Pooltype---Primary-.
>--+-----+-----+-----+-----+-----+-----+----->
    '-Pooltype---Primary-' '-Description---description-'

    .-CloudType---Swift-----.
>--+-----+-----+-----+-----+-----+-----+----->
    '-CloudType---+Azure-----+'
                    +-S3-----+
                    +-IBMCloudswift+
                    +-Swift-----+
                    '-V1Swift-----'

                                                    (1)
>--CLOUDUrl---cloud_url--IDentity---cloud_identity----->
>--PAssword---password----->
    .-CLOUDLocation---Offpremise-----.
>--+-----+-----+-----+-----+-----+-----+----->
    '-CLOUDLocation---+Offpremise--+'
                    '-ONpremise--'

>--+-----+-----+-----+-----+-----+-----+----->
    |                                     (2) |
    '-BUCKETName---bucket_name-----'

    .-ACcEss---READWrite-----.
>--+-----+-----+-----+-----+-----+-----+----->
    '-ACcEss---+READWrite---+'
                    +-READOnly-----+
                    '-UNAVailable-'

    .-MAXWriters---NOLimit-----.
>--+-----+-----+-----+-----+-----+-----+----->
    '-MAXWriters---+NOLimit-----+'
                    '-maximum_writers-'

    .-REUsedelay---1-----. .-ENCRypt---Yes-----.
>--+-----+-----+-----+-----+-----+-----+----->
    '-REUsedelay---days-' |                                     (3) |
                    '-ENCRypt---+Yes+-----+'
                    '-No--'

    .-COMPRession---Yes-----.
>--+-----+-----+-----+-----+-----+-----+-----><
    '-COMPRession---+Yes---+'
                    '-No--'
```

Notes:

1. If you specified CLOUDTYPE=AZURE, do not specify the IDENTITY parameter.
2. This parameter is valid only if you specify CLOUDTYPE=S3.
3. The default value of the ENCRYPT parameter is conditional. The server encrypts data by default if the CLOUDLOCATION parameter is set to OFFPREMISE. If the CLOUDLOCATION parameter is set to ONPREMISE, the default is No.

Parameters

pool_name (Required)

Specifies the cloud-container storage pool to define. This parameter is required. The maximum length of the name is 30 characters.

STGType=Cloud (Required)

Specifies the type of storage that you want to define for a cloud-container storage pool. To ensure that the storage pool can be used in a cloud environment, you must specify STGTYPE=CLOUD.

Tip: To optimize performance, set up one or more local storage directories to temporarily hold data that is moving to the cloud. After you define a cloud-container storage pool, use the DEFINE STGPOOLDIRECTORY command to assign local directories to the cloud-container storage pool.

POoltype=PRimary

Specifies that you want to define a primary storage pool. This parameter is optional.

DEscription

Specifies a description of the cloud-container storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

CLOUDType

Specifies the type of cloud environment where you are configuring the storage pool.

You can specify one of the following values:

Azure

Specifies that the storage pool uses a Microsoft Azure cloud computing system. If you define a storage pool as using Azure with this parameter, you cannot later change the storage pool type by using the UPDATE STGPOOL command.

S3

Specifies that the storage pool uses a cloud computing system with the Simple Storage Service (S3) protocol, such as IBM® Cloud Object Storage or Amazon Web Services (AWS) S3. If you define a storage pool as using S3 with this parameter, you cannot later change the storage pool type by using the UPDATE STGPOOL command.

IBMCloudswift

Specifies that the storage pool uses an IBM Cloud cloud computing system with an OpenStack Swift cloud computing system.

SWift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 2 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol that it is using.

V1Swift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 1 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol that it is using.

This parameter is optional. If you do not specify the parameter, the default value, SWIFT, is used.

CLOUDUrl

Specifies the URL of the cloud environment where you are configuring the storage pool. Based on your cloud provider, you can use a blob service endpoint, region endpoint URL, an accesser IP address, a public authentication endpoint, or a similar value for this parameter. Be sure to include the protocol, such as `https://` or `http://`, at the beginning of the URL. The maximum length of the web address is 870 characters. The CLOUDURL parameter is not validated until the first backup begins.

For more information about how to locate these values, select your cloud service provider from the list on the Configuring a cloud-container storage pool for data storage page.

Tip: To use more than one IBM Cloud Object Storage accesser, list the accesser IP addresses separated by a vertical bar (|), with no spaces, such as in the following example:

```
CLOUDURL=<accesser_URL1>|<accesser_URL2>|<accesser_URL3>
```

If you are using the Operations Center, type an accesser IP address in the URL field of the Add Storage pool wizard, and then press Enter to add additional IP addresses. Use multiple accessers to improve performance.

This parameter is required if you specify the CLOUDTYPE parameter.

- Azure
- S3 (Simple Storage Service)
- IBMCloudswift
- SWift
- V1Swift

Identity

Specifies the user ID for the cloud that is specified in the STGTYPE=CLOUD parameter. This parameter is required for all supported cloud computing systems except Azure. If you specified CLOUDTYPE=AZURE, do not specify the IDENTITY parameter. Based on your cloud provider, you can use an access key ID, a user name, a tenant name and user name, or a similar value for this parameter. The maximum length of the user ID is 255 characters.

PASsword (Required)

Specifies the password for the cloud that is specified in the STGTYPE=CLOUD parameter. Based on your cloud provider, you can use a shared access signature (SAS) token, secret access key, an API key, a password, or a similar value for this parameter. This parameter is required. The maximum length of the password is 255 characters. The IDENTITY and PASSWORD parameters are not validated until the first backup begins.

CLOUDLocation

Specifies the physical location of the cloud that is specified in the CLOUD parameter. This parameter is optional. The default value is OFFPREMISE. You can specify one of the following values:

- OFFpremise
- ONpremise

BUCKETName

Specifies the name for an AWS S3 bucket or a IBM Cloud Object Storage vault to use with this storage pool, instead of using the default bucket name or vault name. This parameter is optional, and is valid only if you specify CLOUDTYPE=S3. If the name that you specify does not exist, the server creates a bucket or vault with the specified name before using the bucket or vault. Follow the naming restrictions for your cloud provider when specifying this parameter. Review the permissions for the bucket or vault and make sure that the credentials for this storage pool have permission to read, write, list, and delete objects in this bucket or vault. If you do not have the ability to change or view the permissions, and you have not already written data to this storage pool, use the UPDATE STGPOOL command with the BUCKETNAME parameter to use a different bucket or vault.

ACCess

Specifies how client nodes and server processes access the cloud-container storage pool. This parameter is optional. The default value is READWRITE. You can specify one of the following values:

READWrite

Specifies that client nodes and server processes can read and write to the cloud-container storage pool. This value is the default.

READOnly

Specifies that client nodes and server processes can read only from the cloud-container storage pool.

UNAVailable

Specifies that client nodes and server processes cannot access the cloud-container storage pool.

MAXWriters

Specifies the maximum number of writing sessions that can run concurrently on the cloud-container storage pool. Specify a maximum number of writing sessions to control the performance of the cloud-container storage pool from negatively impacting other system resources. This parameter is optional. The default value is NOLIMIT. You can specify one of the following values:

NOLimit

Specifies that no maximum size limit exists for the number of writers that you can use. This value is the default.

maximum_writers

Limits the maximum number of writers that you can use. Specify an integer in the range 1 - 99999.

REUsedelay

Specifies the number of days that must elapse after all deduplicated extents are removed from a cloud-container storage pool. This parameter controls the duration that deduplicated extents are associated with a cloud-container storage pool. When the value that is specified for the parameter expires, the deduplicated extents are deleted from the cloud-container storage pool. The default is 1. You can specify one of the following values:

1

Specifies that deduplicated extents are deleted from a cloud-container storage pool after one day. This value is the default.

days

You can specify an integer in the range 0 - 9999.

Tip: Set this parameter to a value that is greater than the number specified for the SET DRMDBBACKUPEXPIREDDAYS command. If you set this parameter to a higher value, you can ensure that when you restore the database to an earlier level, the references to files in the cloud-container storage pool are still valid.

ENCRypt

Specifies whether the server encrypts client data before it writes it to the storage pool. You can specify the following values:

Yes

Specifies that client data is encrypted by the server.

No

Specifies that client data is not encrypted by the server.

This parameter is optional. The default depends on the physical location of the cloud, which is specified by the CLOUDLOCATION parameter. If the cloud is off premise, the server encrypts data by default. If the cloud is on premises, the server does not encrypt data by default.

COMPRession

Specifies whether data is compressed in the storage pool. This parameter is optional. You can specify one of the following values:

No

Specifies that data is not compressed in the storage pool.

Yes

Specifies that data is compressed in the storage pool. This is the default.

Example 1: Define an OpenStack Swift cloud-container storage pool

Define an OpenStack Swift cloud-container storage pool that is named STGPOOL1.

```
define stgpool stgpool1 stgtype=cloud
cloudtype=swift cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 description="OpenStack Swift cloud"
```

Example 2: Define a cloud-container primary storage pool

Define a cloud-container primary storage pool that is named STGPOOL1.

```
define stgpool stgpool1 stgtype=cloud
cloudtype=swift cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 pooltype=primary
```

Example 3: Define a cloud-container storage pool with read only access

Define a cloud-container storage pool that is named STGPOOL1 with read only access.

```
define stgpool stgpool1 stgtype=cloud
cloudtype=swift cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 access=readonly
```

Example 4: Define a cloud-container storage pool with 99 writing sessions

Define a cloud-container storage pool that is named STGPOOL1 with 99 writing sessions.

```
define stgpool stgpool1 stgtype=cloud
cloudtype=swift cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 maxwr=99
```

Example 5: Define a cloud-container storage pool in which deduplicated extents are deleted after two days

Define a cloud-container storage pool that is named STGPOOL1 and deduplicated extents are deleted after two days.

```
define stgpool stgpool1 stgtype=cloud
cloudtype=swift cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 reusedelay=2
```

Related tasks:

Configuring a cloud-container storage pool for data storage

Related information:

DEFINE STGPOOL (Define a directory-container storage pool)

Use this command to define a directory-container storage pool that is used for data deduplication.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--STGType--===-Directory----->
. -Pooltype--===-Primary-.
>--+-----+-----+-----+-----+-----+-----+----->
' -Pooltype--===-Primary-' '-DEscription--===-description-'
. -ACcess--===-READWrite-----.
>--+-----+-----+-----+-----+-----+-----+----->
' -ACcess--===-+READWrite---+-'
      +READOnly---+
      '-UNAVailable-'
. -MAXSize--===-NOLimit------.
>--+-----+-----+-----+-----+-----+-----+----->
' -MAXSize--===-+NOLimit-----+-'
      '-maximum_file_size-'
. -MAXWriters--===-NOLimit------.
>--+-----+-----+-----+-----+-----+-----+----->
' -MAXWriters--===-+NOLimit-----+-'
      '-maximum_writers-'
>--+-----+-----+-----+-----+-----+-----+----->
' -NEXTstgpool--===-pool_name-'
>--+-----+-----+-----+-----+-----+-----+----->
' -PROTECTstgpool--===-target_stgpool-'
>--+-----+-----+-----+-----+-----+-----+----->
|                                     .,-----, |
|                                     V           ||
' -PROTECTLOCalstgpool--===-local_target_stgpool--+'
. -REUsedelay--===-1----. . -ENCRypt--===-No-----.
>--+-----+-----+-----+-----+-----+-----+----->
' -REUsedelay--===-days-' '-ENCRypt--===-+Yes+-'
                                     '-No--'
. -COMPRession--===-Yes-----.
>--+-----+-----+-----+-----+-----+-----+----->
' -COMPRession--===-+Yes+-'
                                     '-No--'
```

Parameters

pool_name (Required)

Specifies the storage pool to define. This parameter is required. The maximum length of the name is 30 characters.

STGType=Directory (Required)

Specifies the type of storage that you want to define for a storage pool. This parameter specifies that a directory-container type of storage pool is assigned to the storage pool. You must define a storage pool directory for this type of storage pool by using the DEFINE STGPOOLDIRECTORY command.

Requirements:

- Ensure that enough space is available on the file system for the directory-container storage pool.

- You must store the directory-container storage pool and the DB2® database on separate mount points on the file system. The directory-container storage pool might grow to occupy all the space on the directory it is stored on.
- You must use a file system other than the file system where the IBM Spectrum Protect™ server is located.

POoltype=Primary

Specifies that you want the storage pool to be used as a primary storage pool. This parameter is optional.

DESCription

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

ACCess

Specifies how client nodes and server processes can access the storage pool. This parameter is optional. You can specify one of the following values:

READWrite

Specifies that client nodes and server processes can read and write to the storage pool.

READOnly

Specifies that client nodes and server processes can read only from the storage pool.

UNAVailable

Specifies that client nodes and server processes cannot access the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is NOLIMIT. You can specify one of the following values:

NOLimit

Specifies that there is no maximum size limit for physical files that are stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer in the range 1 - 999999, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 GB. You can use one of the following scale factors:

Table 1. Scale factor for the maximum file size

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

Tip: If you do not specify a unit of measurement for the maximum file size, the value is specified in bytes.

When the physical size of the storage pool exceeds the MAXSIZE parameter, the following table shows where files are typically stored.

Table 2. The location of a file according to the file size and the pool that is specified

Pool that is specified	Result
No pool is specified as the next storage pool in the hierarchy.	The server does not store the file.
A pool is specified as the next storage pool in the hierarchy.	The server stores the file in the storage pool that you specified.

Tip: If you also specify the NEXTstgpool parameter, define one storage pool in your hierarchy to have no limit on the maximum file size by specifying the MAXSIZE=NOLimit parameter. When you have at least one pool with no size limit, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent during data deduplication processing, the server considers the size of the data deduplication process to be the file size. If the total size of all files in the process is larger than the maximum size limit, the server does not store the files in the storage pool.

MAXWriters

Specifies the maximum number of I/O threads for the following processes:

- The number of I/O threads that can run concurrently on the directory-container storage pool.
- The number of I/O threads that are written simultaneously to the directory-container storage pool.

This parameter is optional. As a best practice, use the default value of NOLIMIT. You can specify the following values:

NOLimit

Specifies that no maximum number of I/O threads are written to the storage pool.

maximum_writers

Limits the maximum number of I/O threads that you can use. Specify an integer in the range 1 - 99999.

Tip: The IBM Spectrum Protect server manages the number of I/O threads automatically based on the resources that are available and the server load.

NEXTstgpool

Specifies the name of a random-access or primary sequential storage pool to which files are stored when the directory-container storage pool is full. This parameter is optional.

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.
- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.

PROTECTstgpool

Specifies the name of the directory-container storage pool on the target replication server where the data is backed up when you use the PROTECT STGPOOL command for this storage pool. This parameter is optional.

PROTECTLOCstgpools

Specifies the name of the container-copy storage pool on a local device where the data is backed up. This container-copy storage pool will be a local target storage pool when you use the PROTECT STGPOOL command. You can specify a maximum of two container-copy storage pool names. Separate multiple names with commas and no intervening spaces. The maximum length of each name is 30 characters. This parameter is optional.

REUsedelay

Specifies the number of days that must elapse before all deduplicated extents are removed from a directory-container storage pool. This parameter controls the duration that deduplicated extents are associated with a directory-container storage pool after they are no longer referenced. When the value that is specified for the parameter expires, the deduplicated extents are deleted from the directory-container storage pool. Specify an integer in the range 0 - 9999. The default value for directory-container storage pools is 1, which means that deduplicated extents that are no longer referenced are deleted from a directory-container storage pool after 1 day.

Set this parameter to a value greater than the number that is specified as your database backup period to ensure that data extents are still valid when you restore the database to another level.

ENCRypt

Specifies whether the server encrypts client data before the server writes the data to the storage pool. You can specify the following values:

Yes

Specifies that client data is encrypted by the server.

No

Specifies that client data is not encrypted by the server. This is the default value.

COMPReSSion

Specifies whether data is compressed in the storage pool. This parameter is optional. You can specify one of the following values:

No

Specifies that data is not compressed in the storage pool.

Yes

Specifies that data is compressed in the storage pool. This is the default.

Example: Define a directory-container storage pool that is configured for overflow storage when the storage pool is full

Define a directory-container storage pool that is named STGPOOL1. The storage pool is configured for overflow storage to a tape storage pool when the storage pool is full.

```
define stgpool stgpool1 stgtype=directory nextstgpool=overflow_tape_pool
```

Example: Define a directory-container storage pool that specifies the maximum file size

Define a directory-container storage pool that is named STGPOOL2. The storage pool specifies the maximum file size that the server can store in the storage pool as 100 megabytes.

```
define stgpool stgpool2 stgtype=directory maxsize=100M
```

Example: Define a directory-container storage pool on the source replication server with a directory-container storage pool on the target replication server to back up data

Define a directory-container storage pool that is named STGPOOL3. The data for storage pool STGPOOL3 is backed up to a directory-container storage pool, TARGET_STGPOOL3 on the target replication server.

```
define stgpool stgpool3 stgtype=directory protectstgpool=target_stgpool3
```

Example: Define a directory-container storage pool on the source replication server with a container-copy storage pool to back up data locally

Define a directory-container storage pool that is named STGPOOL3. The data for storage pool STGPOOL3 is backed up to a local container-copy storage pool, TARGET_LOCALSTGPOOL.

```
define stgpool stgpool3 stgtype=directory protectlocalstgpool=target_localstgpool
```

Example: Define a directory-container storage pool and disable compression

Define a directory-container storage pool that is named STGPOOL1 and disable compression.

```
define stgpool stgpool1 stgtype=directory compression=no
```

Table 3. Commands related to DEFINE STGPOOL (Define a directory-container storage pool)

Command	Description
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
QUERY CONTAINER	Displays information about a container.
QUERY STGPOOL	Displays information about storage pools.
REPAIR STGPOOL	Repairs a directory-container storage pool.
UPDATE STGPOOL (directory-container)	Update a directory-container storage pool.

DEFINE STGPOOL (Define a container-copy storage pool)

Use this command to define a container-copy storage pool to hold a copy of data from a directory-container storage pool.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--device_class_name----->
```

```
>--Pooltype----COPYContainer--MAXScratch----number----->
```



```

>----->
'-DEscription---description-'

.-ACCess---READWrite-----
>----->
'-ACCess---+READWrite---+'
      +-READOnly----+
      '-UNAVailable-'

.-PROTECTProcess---2----- .-REClaim---100-----
>----->
'-PROTECTProcess---number-' '-REClaim---percent-'

.-RECLAIMLiMit---NOLimit-----
>----->
'-RECLAIMLiMit---+NOLimit---+'
      '-vol_limit-'

.-REUsedelay---0-----
>----->
'-REUsedelay---days-'

```

Parameters

pool_name (Required)

Specifies the name of the container-copy storage pool. The name must be unique, and the maximum length is 30 characters.

device_class_name (Required)

Specifies the name of the sequential access device class to which this storage pool is assigned.

Restriction: You cannot specify the following device class types:

- DISK
- FILE
- CENTERA
- NAS
- REMOVABLEFILE
- SERVER

Restriction: Virtual tape libraries are not supported, regardless of which library type is defined. Only physical tape is supported.

POoltype=COPYCONtainer (Required)

Specifies that you want to define a container-copy storage pool. A container-copy storage pool is used only to store a copy of data from a directory-container storage pool.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer in the range 0 - 100000000. If the server can request scratch volumes as needed, you do not have to define each volume to be used.

The value of this parameter is used to estimate the total number of volumes that are available in the storage pool and the corresponding estimated capacity for the storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the storage pool until the access mode is changed. An administrator can then query the server for empty, offsite scratch volumes and return them to the onsite location.

DEscription

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCess

Specifies how server processes such as storage-pool protection and repair can access data in the storage pool. This parameter is optional. The default value is READWRITE. You can specify one of the following values:

READWrite

Specifies that the server can read and write to volumes in the storage pool.

READOnly

Specifies that the server can only read volumes in the storage pool. The server can use data in the storage pool to restore extents to directory-container storage pools. No operations that write to the container-copy storage pool are allowed.

UNAVailable

Specifies that the server cannot access data that is stored on volumes in the storage pool.

PROTECTProcess

Specifies the maximum number of parallel processes that are used when you issue the PROTECT STGPOOL command to copy data to this pool from a directory-container storage pool. This parameter is optional. Enter a value in the range 1 - 20. The default value is 2.

The time that is required to complete the copy operation might be decreased by using multiple, parallel processes. However, in some cases when multiple processes are running, one or more of the processes must wait to use a volume that is already in use by a different process.

When you specify this value, consider the number of logical and physical drives that can be dedicated to the copy operation. To access a tape volume, the server uses a mount point and a drive. The number of available mount points and drives depends on the mount limit of the device class for the storage pool, and on other server and system activity.

This parameter is ignored if you use the PREVIEW=YES option on the PROTECT STGPOOL command. In that case, only one process is used and no mount points or drives are needed.

REClaim

Specifies when a volume becomes eligible for reclamation and reuse. Specify eligibility as the percentage of a volume's space that is occupied by extents that are no longer stored in the associated directory-container storage pool. Reclamation moves any extents that are still stored in the associated directory-container storage pool from eligible volumes to other volumes. Reclamation occurs only when a PROTECT STGPOOL command stores data into this storage pool.

This parameter is optional. You can specify an integer in the range 1 - 100. The default value is 100, which means that volumes in this storage pool are not reclaimed.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

By setting the reclaim value to 50 percent or greater, data that is moved from two reclaimed volumes uses no more than the equivalent of one new volume.

Use caution when you use reclamation with container-copy storage pools that have offsite volumes. When an offsite volume becomes eligible for reclamation, in effect the server moves the extents on the volume back to the onsite location. If a disaster occurs onsite, the server can obtain extents from the offsite volume if the restored database refers to extents on the offsite volume. Therefore, for disaster recovery purposes, ensure that you schedule database backups to run after storage pool protection schedules and DRM move schedules have run, and ensure that all database backup volumes are taken offsite along with the DRM volumes.

Tip: Set different reclamation values for offsite container-copy storage pools and onsite container-copy storage pools. Because container-copy storage pools store deduplicated data, the data extents are spread across multiple tape volumes. When you choose a reclamation threshold for an offsite copy, carefully consider the number of available mount points and the number of tape volumes that you must retrieve if a disaster occurs. Setting a higher threshold means that you must retrieve more volumes than you would if your reclamation value was lower. Using a lower threshold reduces the number of mount points that are required in a disaster. The preferred method is to set the reclamation value for offsite copies to 60, and for onsite copies, in the range 90 - 100.

RECLAIMLimit

Specifies the maximum number of volumes that the server reclaims when you issue the PROTECT STGPOOL command and specify the RECLAIM=YESLIMITED or RECLAIM=ONLYLIMITED option. This parameter is valid only for container-copy storage pools. This parameter is optional. The default value is NOLIMIT. You can specify one of the following values:

NOLimit

Specifies that all volumes in the container-copy storage pool are processed for reclamation.

vol_limit

Specifies the maximum number of volumes in the container-copy storage pool that are reclaimed. The value that you specify determines how many new scratch tapes are available after reclamation processing completes. You can specify a number in the range 1 - 100000.

REUsedelay

Specifies the number of days that must elapse after all extents are deleted from a volume before the volume can be rewritten or returned to scratch status. This parameter is optional. You can specify an integer in the range 0 - 9999. The default value is 0, which means that a volume can be rewritten or returned to scratch status as soon as all the extents are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to extents in the storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. If you use disaster recovery manager, the number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDDAYS command.

Example: Define a container-copy storage pool with an LTO7A device class

Define a container-copy storage pool, CONTAINER1_COPY2, to the LTO7A device class. Allow up to 50 scratch volumes for this pool. Delay the reuse of volumes for 45 days.

```
define stgpool container1_copy2 lto7a pooltype=copycontainer
maxscratch=50 reusedelay=45
```

Table 1. Commands related to DEFINE STGPOOL (Define a container-copy storage pool)

Command	Description
DEFINE STGPOOL (directory-container)	Define a directory-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
QUERY STGPOOL	Displays information about storage pools.
REPAIR STGPOOL	Repairs a directory-container storage pool.
UPDATE STGPOOL (container-copy)	Update a container-copy storage pool that stores copies of data from a directory-container storage pool.
UPDATE STGPOOL (directory-container)	Update a directory-container storage pool.

DEFINE STGPOOL (Define a primary storage pool assigned to random access devices)

Use this command to define a primary storage pool that is assigned to random access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-DEFine STGpool--pool_name--DISK-----+-----+----->
                                     .-Pooltype----Primary-.
                                     '-Pooltype----Primary-'

    .-STGType----Devclass-.
>--+-----+-----+-----+-----+-----+----->
    '-STGType----Devclass-' '-DESCRIPTION----description-'

    .-ACCESS----READWrite-----
>--+-----+-----+-----+-----+-----+----->
    '-ACCESS----+READWrite----+'
                                     +READOnly----+
                                     '-UNAVailable-'

    .-MAXSize----NOLimit----- .-CRCDData----No-----
>--+-----+-----+-----+-----+-----+----->
    '-MAXSize----maximum_file_size-' '-CRCDData----+Yes+-'
                                     '-No--'

                                     .-Highmig----90-----
>--+-----+-----+-----+-----+-----+----->
    '-NEXTstgpool----pool_name-' '-Highmig----percent-'

```


Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is NOLIMIT. You can specify the following values:

NOLimit

Specifies that there is no maximum size limit for physical files that are stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer 1 - 999999 terabytes, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 GB. You can use one of the following scale factors:

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

The client estimates the size of files that are sent to the server. The client estimate is used rather than the actual amount of data that is sent to the server. Client options, such as deduplication, compression, and encryption, can cause the actual amount of data that is sent to the server to be larger or smaller than the size estimate. For example, the compression of a file might be smaller in size than the estimate, thus sending less data than the estimate. Furthermore, a binary file might be larger in size after the compression processing, thus sending more data than the estimate.

When the physical size of the storage pool exceeds the MAXSIZE parameter, the following table shows where files are typically stored.

Table 1. The location of a file according to the file size and the pool that is specified

File size	Pool specified	Result
Exceeds the maximum size	No pool is specified as the next storage pool in the hierarchy	The server does not store the file
	A pool is specified as the next storage pool in the hierarchy	The server stores the file in the next storage pool that can accept the file size

Tip: If you also specify the NEXTstgpool parameter, define one storage pool in your hierarchy to have no limit on the maximum file size by specifying the MAXSize=NOLimit parameter. When you have at least one pool with no size limit, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent in a single transaction, the server considers the size of the transaction to be the file size. If the total size of all files in the transaction is larger than the maximum size limit, the server does not store the files in the storage pool.

CRCDData

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more expenditure is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

NEXTstgpool

Specifies a primary storage pool to which files are migrated. This parameter is optional.

If you do not specify a next storage pool, the following actions occur:

- The server cannot migrate files from this storage pool
- The server cannot store files that exceed the maximum size for this storage pool in another storage pool

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.
- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.

HIghmig

Specifies that the server starts migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer 0 - 100. The default value is 90.

When the storage pool exceeds the high migration threshold, the server can start migration of files by node, to the next storage pool. The NEXTSTGPOOL parameter defines this setting. You can specify HIGHMIG=100 to prevent migration for this storage pool.

LOWmig

Specifies that the server stops migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer 0 - 99. The default value is 70.

When migration is by node or file space, depending upon collocation, the level of the storage pool can fall below the value that you specified for this parameter. To empty the storage pool, set LOWMIG=0.

CAChe

Specifies whether the migration process leaves a cached copy of a file in this storage pool after you migrate the file to the next storage pool. This parameter is optional. The default value is NO. You can specify the following values:

Yes

Specifies that caching is enabled.

No

Specifies that caching is disabled.

Using cache might improve the ability to retrieve files, but might affect the performance of other processes.

MIGPRocess

Specifies the number of processes that the server uses for migrating files from this storage pool. This parameter is optional. You can specify an integer 1 - 999. The default value is 1.

During migration, these processes are run in parallel to provide the potential for improved migration rates.

Tips:

- The number of migration processes is dependent upon the following settings:
 - The MIGPROCESS parameter
 - The collocation setting of the next pool
 - The number of nodes or the number of collocation groups with data in the storage pool that is being migrated

For example, suppose that `MIGPROCESS =6`, the next pool `COLLOCATE` parameter is set to `NODE`, but there are only two nodes with data on the storage pool. Migration processing consists of only two processes, not six. If the `COLLOCATE` parameter is set to `GROUP` and both nodes are in the same group, migration processing consists of only one process. If the `COLLOCATE` parameter is set to `NO` or `FILESPEC`, and each node has two file spaces with backup data, then migration processing consists of four processes.

- When you specify this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool that is defined to the target storage pool.

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. To calculate a value to compare to the specified `MIGDELAY` value, the server counts the following items:

- The number of days that the file was in the storage pool
- The number of days, if any, since the file was retrieved by a client

The lesser of the two values are compared to the specified `MIGDELAY` value. For example, if all the following conditions are true, a file is not migrated:

- A file was in a storage pool for five days.
- The file was accessed by a client within the past three days.
- The value that is specified for the `MIGDELAY` parameter is four days.

This parameter is optional. You can specify an integer 0 - 9999. The default is 0, which means that you do not want to delay migration.

If you want the server to count the number of days that are based on when a file was stored and not when it was retrieved, use the `NORETRIEVEDATE` server option.

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional. The default is `YES`.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

You can specify one of the following values:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

No

Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

AUTOCopy

Specifies when IBM Spectrum Protect™ runs simultaneous-write operations. The default value is `CLIENT`. This parameter is optional and affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These

pools remain active for the duration of the migration process. Copy storage pools are specified using the COPYSTGPOOLS parameter. Active-data pools are specified using the ACTIVEDATAPOOLS parameter.

You can specify one of the following values:

None

Specifies that the simultaneous-write function is disabled.

Client

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

MIGRation

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

All

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

COPYSTGpools

Specifies the names of copy storage pools where the server simultaneously writes data. The COPYSTGPOOLS parameter is optional. You can specify a maximum of three copy pool names that are separated by commas. Spaces between the names of the copy pools are not allowed. When you specify a value for the COPYSTGPOOLS parameter, you can also specify a value for the COPYCONTINUE parameter.

The combined total number of storage pools that are specified in the COPYSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the COPYCONTINUE value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to copy storage pools during the following operations:

- Back up and archive operations by IBM Spectrum Protect backup-archive clients or application clients that are using the IBM Spectrum Protect API
- Migration operations by IBM Spectrum Protect for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

Restriction: The simultaneous-write function is not supported for the following store operations:

- When the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is followed.
- NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the management class has copy storage pools that are defined:
 - The copy storage pools are ignored
 - The data is stored into the primary storage pool only

Attention: The function that is provided by the COPYSTGPOOLS parameter is not intended to replace the BACKUP STGPOOL command. If you use the COPYSTGPOOLS parameter, continue to use the BACKUP STGPOOL command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

COPYContinue

Specifies how the server usually reacts to a copy storage pool write failure for any of the copy storage pools that are listed in the COPYSTGPOOLS parameter. This parameter is optional. The default value is YES. When you specify the COPYCONTINUE parameter, you must also specify the COPYSTGPOOLS parameter.

You can specify the following values:

Yes

If the COPYCONTINUE parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

No

If the COPYCONTINUE parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the COPYCONTINUE parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The ACTIVEDATAPOOLS parameter is optional. Spaces between the names of the active-data pools are not allowed.

The combined total number of storage pools that are specified in the COPYSGTPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool that is specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to active-data pools only during backup operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API.

Restrictions:

1. This parameter is available only to primary storage pools that use "NATIVE" or "NONBLOCK" data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Writing data simultaneously to active-data pools is not supported when you use LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is followed.
3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the TOCDESTINATION in the copy group of the management class has active-data pools that are defined:
 - The active-data pools are ignored
 - The data is stored into the primary storage pool only
4. You cannot use the simultaneous-write function with CENTERA storage devices.
5. Data that is being imported is not stored in active-data pools. After an import operation, use the COPY ACTIVEDATA command to store the imported data in an active-data pool.

Attention: The function that is provided by the ACTIVEDATAPOOLS parameter is not intended to replace the COPY ACTIVEDATA command. If you use the ACTIVEDATAPOOLS parameter, use the COPY ACTIVEDATA command to ensure that the active-data pools contain all active data of the primary storage pool.

SHRED

Specifies whether data is physically overwritten when it is deleted. This parameter is optional. You can specify an integer 0 - 10. The default value is 0.

If you specify a value of zero, the server deletes the data from the database. However, the storage that is used to contain the data is not overwritten, and the data exists in storage until that storage is reused for other data. It might be possible to

discover and reconstruct the data after it is deleted.

If you specify a value greater than zero, the server deletes the data both logically and physically. The server overwrites the storage that is used to contain the data the specified number of times. This overwriting increases the difficulty of discovering and reconstructing the data after it is deleted.

To ensure that all copies of the data are shredded, specify a SHRED value greater than zero for the storage pool that is specified in the NEXTSTGPOOL parameter. Do not specify either the COPYSTGPOOLS or ACTIVEPOOL. Specifying relatively high values for the overwrite count generally improves the level of security, but might affect performance adversely.

Overwriting of deleted data is done asynchronously after the delete operation is complete. Therefore, the space that is occupied by the deleted data remains occupied for some time. The space is not available as free space for new data.

A SHRED value greater than zero cannot be used if the value of the CACHE parameter is YES.

Important: After an export operation finishes and identifies files for export, any change to the storage pool SHRED value is ignored. An export operation that is suspended retains the original SHRED value throughout the operation. You might want to consider canceling your export operation if changes to the storage pool SHRED value jeopardize the operation. You can reissue the export command after any needed cleanup.

Example: Define a primary storage pool for a DISK device class

Define a primary storage pool, POOL1, to use the DISK device class, with caching enabled. Limit the maximum file size to 5 MB. Store any files larger than 5 MB in subordinate storage pools that begin with the PROG2 storage pool. Set the high migration threshold to 70 percent, and the low migration threshold to 30 percent.

```
define stgpool pool1 disk
description="main disk storage pool" maxsize=5m
highmig=70 lowmig=30 cache=yes
nextstgpool=prog2
```

DEFINE STGPOOL (Define a primary storage pool assigned to sequential access devices)

Use this command to define a primary storage pool that is assigned to sequential access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--device_class_name----->
. -Pooltype---PRimary-. .-STGType---Devclass-.
>-----+-----+----->
'-Pooltype---PRimary-' '-STGType---Devclass-'
>-----+-----+----->
'-DESCRiption---description-'
. -ACCess---READWrite-----
>-----+-----+----->
'-ACCess---+READWrite---+
+READOnly---+
'-UNAVailable-'
. -MAXSize---NOLimit-----
>-----+-----+----->
| (1) (2) |
'-MAXSize---maximum_file_size-----'
. -CRCDATA---No-----
>-----+-----+----->
```



```

'-DATAFormat-----+NATive-----+-----'
      +-NONblock----+
      +-NETAPPDump--+
      +-CELERRADump--+
      '-NDMPDump----'

.-AUTOCopy----CLient-----
>-----+-----+-----+----->
'-AUTOCopy----+None-----+'
      +-CLient----+
      +-MIGRation+
      '-All-----'

>-----+-----+-----+----->
|                                     |
|                                     |
|          v                          | (1) (2) | |
|'-COPYSTGpools-----copy_pool_name-----+'

.-COPYContinue----Yes-----
>-----+-----+-----+----->
|                                     |
|                                     | (1) (2) |
|'-COPYContinue----+Yes-----+'
|                                     |
|                                     | '-No--'

>-----+-----+-----+----->
|                                     |
|                                     |
|          v                          | | |
|'-ACTIVEDATApools-----active-data_pool_name-----+'

.-DEDuplicate----No-----
>-----+-----+-----+----->
'-DEDuplicate----+No-----+'
|                                     | (5) |
|                                     | '-Yes-----'

.-IDENTIFYProcess----1-----
>-----+-----+-----+----->>
|                                     | (6) |
|'-IDENTIFYProcess----number-----'

```

Notes:

1. This parameter is not available for storage pools that use the data formats NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
2. This parameter is not available or is ignored for CENTERA storage pools.
3. The RECLAMATIONTYPE=SNAPLOCK setting is valid only for storage pools that are defined to servers that are enabled for IBM Spectrum Protect™ for Data Retention. The storage pool must be assigned to a FILE device class, and the directories that are specified in the device class must be NetApp SnapLock volumes.
4. The values NETAPPDUMP, CELERRADUMP, and NDMPDUMP are not valid for storage pools that are defined with a FILE-type device class.
5. This parameter is valid only for storage pools that are defined with a FILE-type device class.
6. This parameter is available only when the value of the DEDuplicate parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be defined. The name must be unique, and the maximum length is 30 characters.

device_class_name (Required)

Specifies the name of the device class to which this storage pool is assigned. You can specify any device class except for the DISK device class.

Pooltype=Primary

Specifies that you want to define a primary storage pool. This parameter is optional. The default value is PRIMARY.

STGType

Specifies the type of storage that you want to define for a storage pool. This parameter is optional. The default value is DEVCLASS.

Devclass

Specifies that a device class type of storage pool is assigned to the storage pool.

DESCRiption

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCess

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. The default value is READWRITE. You can specify the following values:

READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSIze

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is NOLIMIT. You can specify one of the following values:

NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer from 1 to 999999 terabytes, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

The client estimates the size of files that are sent to the server. The client estimate is used rather than the actual amount of data that is sent to the server. Client options, such as deduplication, compression, and encryption, can cause the actual amount of data that is sent to the server to be larger or smaller than the size estimate. For example, the compression of a file might be smaller in size than the estimate, thus sending less data than the estimate. Furthermore, a binary file might be larger in size after the compression processing, thus sending more data than the estimate.

When the physical size of the storage pool exceeds the MAXSIZE parameter, the following table shows where files are typically stored.

Table 1. The location of a file according to the file size and the pool that is specified

File size	Pool specified	Result
Exceeds the maximum size	No pool is specified as the next storage pool in the hierarchy	The server does not store the file
	A pool is specified as the next storage pool in the hierarchy	The server stores the file in the next storage pool that can accept the file size

Tip: If you also specify the NEXTstgpool parameter, define one storage pool in your hierarchy to have no limit on the maximum file size by specifying the MAXSIze=NOLimit parameter. When you have at least one pool with no size limit, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent in a single transaction, the server considers the size of the transaction to be the file size. If the total size of all files in the transaction is larger than the maximum size limit, the server does not store the files in the storage pool.

Restriction:

This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

CRCDData

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

NEXTstgpool

Specifies a primary storage pool to which files are migrated. You cannot migrate data from a sequential access storage pool to a random access storage pool. This parameter is optional.

If this storage pool does not have a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

When there is insufficient space available in the current storage pool, the NEXTSTGPOOL parameter for sequential access storage pools does not allow data to be stored into the next pool. In this case, the server issues a message and the transaction fails.

For next storage pools with a device type of FILE, the server completes a preliminary check to determine whether sufficient space is available. If space is not available, the server skips to the next storage pool in the hierarchy. If space is available, the server attempts to store data in that pool. However, it is possible that the storage operation might fail because, at the time the actual storage operation is attempted, the space is no longer available.

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.

- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.
- This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP

HIghmig

Specifies that the server starts migration when storage pool utilization reaches this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer 0 - 100. The default value is 90.

When the storage pool exceeds the high migration threshold, the server can start migration of files by volume to the next storage pool defined for the pool. You can set the high migration threshold to 100 to prevent migration for the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

LOwmig

Specifies that the server stops migration when storage pool utilization is at or below this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer 0 - 99. The default value is 70.

When the storage pool reaches the low migration threshold, the server does not start migration of files from another volume. You can set the low migration threshold to 0 to allow migration to empty the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

REClaim

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100. The default value is 60, except for storage pools that use WORM devices.

AIX | **Windows** For storage pools that use a WORM device class, you can lower the value from the default of 100. Lowering the value allows the server to consolidate data onto fewer volumes when needed. Volumes that are emptied by reclamation can be checked out of the library, freeing slots for new volumes. Because the volumes are write-once, the volumes cannot be reused.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP

- NDMPDUMP

RECLAIMPRocess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999. The default value is 1. You can specify one or more reclamation processes for each primary sequential-access storage pool.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Assuming that the RECLAIMSTGPOOL parameter is not specified or that the reclaim storage pool has the same device class as the storage pool that is being reclaimed, each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIMSTGpool

Specifies another primary storage pool as a target for reclaimed data from this storage pool. This parameter is optional. When the server reclaims volumes for the storage pool, the server moves unexpired data from the volumes that are being reclaimed to the storage pool named with this parameter.

A reclaim storage pool is most useful for a storage pool that has only one drive in its library. When you specify this parameter, the server moves all data from reclaimed volumes to the reclaim storage pool regardless of the number of drives in the library.

To move data from the reclaim storage pool back to the original storage pool, use the storage pool hierarchy. Specify the original storage pool as the next storage pool for the reclaim storage pool.

Restriction:

- This parameter is not available for storage pools that use the following data formats:
- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. The following are possible values:

THRESHold

Specifies that volumes that belong to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes that belong to this storage pool are managed for retention using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools that are defined to a server that has data retention protection enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, which is derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume by using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date expires, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools that are being defined must be the same when defined to the same device class name. The DEFINE command can fail if the RECLAMATIONTYPE parameter specified is different from what is defined for storage pools that are already defined to the device class name.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

COLlocate

Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional. The default value is GROUP.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required. Collocation can also impact the number of processes migrating disks to sequential pool.

You can specify one of the following options:

No

Specifies that collocation is disabled. During migration from disk, processes are created at a file space level.

GROup

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.
- During migration from disk, the server creates migration processes at the collocation group level for grouped nodes, and at the node level for ungrouped nodes.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces that are named A, B, C, D, and E. File spaces A and B belong to a file space collocation group but C, D, and E do not. File spaces A and B are collocated by file space collocation group, while C, D, and E are collocated by node.
- During migration from disk, the server creates migration processes at the collocation group level for grouped file spaces.

Data is collocated on the least number of sequential access volumes.

NODe

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

For COLLOCATE=NODE, the server creates processes at the node level when you migrate data from disk.

Filespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

For COLLOCATE=FILESPACE, the server creates processes at the file space level when you migrate data from disk.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer 0 - 100000000. By allowing the server to request scratch volumes, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the storage pool and the corresponding estimated capacity for the storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. When scratch volumes with the device type of FILE are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to files in the storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. All files on a volume must be eligible for migration before the server selects the volume for migration. To calculate a value to compare to the specified MIGDELAY, the server counts the number of days that the file has been in the storage pool.

This parameter is optional. You can specify an integer 0 - 9999. The default is 0, which means that you do not want to delay migration. If you want the server to count the number of days that are based only on when a file was stored and not when it was retrieved, use the NORETRIEVEDATE server option.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional. The default is YES.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

You can specify one of the following values:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

No

Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

MIGPProcess

Specifies the number of parallel processes to use for migrating the files from the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999. The default value is 1.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the migration, and the number of logical and physical drives that can be dedicated to the operation. To access a sequential-access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the migration.

For example, suppose you want to simultaneously migrate the files from volumes in two primary sequential storage pools and that you want to specify three processes for each of the storage pools. The storage pools have the same device class. Assuming that the storage pool to which files are being migrated has the same device class as the storage pool from which files are being migrated, each process requires two mount points and, if the device type is not FILE, two drives. (One drive is for the input volume, and the other drive is for the output volume.) To run six migration processes simultaneously, you need a total of at least 12 mount points and 12 drives. The device class for the storage pools must have a mount limit of at least 12.

If the number of migration processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the migration processes will end. For information about specifying the MOUNTWAIT time, see DEFINE DEVCLASS (Define a device class).

The IBM Spectrum Protect server will start the specified number of migration processes regardless of the number of volumes that are eligible for migration. For example, if you specify ten migration processes and only six volumes are eligible for migration, the server will start ten processes and four of them will complete without processing a volume.

Tip: When you specify this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool that is defined to the target storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

DATAFormat

Specifies the data format to use to back up files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. You can specify the following values:

NATive

Specifies the data format is the native IBM Spectrum Protect server format and includes block headers.

NONblock

Specifies the data format is the native IBM Spectrum Protect server format and does not include block headers. The default minimum block size on a volume that is associated with a FILE device class is 256 KB, regardless how much data is written to the volume. For certain tasks, you can minimize wasted space on storage volumes by specifying the NONBLOCK data format. For example, you can specify the NONBLOCK data format for the following tasks:

- Using content-management products
- Using the DIRMC client option to store directory information
- Migrating very small files by using IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows

In most situations, however, the NATIVE format is preferred.

NETAPPDump

Specifies the data is in a NetApp dump format. This data format must be specified for file system images that are in a dump format and that were backed up from a NetApp or an IBM System Storage® N Series file server that uses NDMP. The server does not complete migration, reclamation, or AUDIT VOLUME for a storage pool with DATAFORMAT=NETAPPDUMP. You can use the MOVE DATA command to move data from one primary storage pool to another, or out of a volume if the volume must be reused.

CELERRADump

Specifies that the data is in an EMC Celerra dump format. This data format must be specified for file system images that are in a dump format and that were backed up from an EMC Celerra file server that uses NDMP. The server does not complete migration, reclamation, or AUDIT VOLUME for a storage pool with DATAFORMAT=CELERRADUMP. You can use the MOVE DATA command to move data from one primary storage pool to another, or out of a volume if the volume must be reused.

NDMPDump

Specifies that the data is in NAS vendor-specific backup format. Use this data format for file system images that were backed up from a NAS file server other than a NetApp or EMC Celerra file server. The server does not complete migration, reclamation, or AUDIT VOLUME for a storage pool with DATAFORMAT=NDMPDUMP. You can use the MOVE DATA command to move data from one primary storage pool to another, or out of a volume if the volume must be reused.

AUTOCopy

Specifies when IBM Spectrum Protect completes simultaneous-write operations. The default value is CLIENT. This parameter is optional and affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If the AUTOCOPY option is set to ALL or CLIENT, and there is at least one storage pool that is listed in the COPYSTGPOOLS or ACTIVEDATAPOOLS options, any client-side deduplication is disabled.

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the COPYSTGPOOLS parameter. Active-data pools are specified using the ACTIVEDATAPOOLS parameter.

You can specify one of the following values:

None

Specifies that the simultaneous-write function is disabled.

CLient

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

MIGRation

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

All

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

COPYSTGPools

Specifies the names of copy storage pools where the server simultaneously writes data. The COPYSTGPOOLS parameter is optional. You can specify a maximum of three copy pool names that are separated by commas. Spaces between the names of the copy pools are not allowed. When you specify a value for the COPYSTGPOOLS parameter, you can also specify a value for the COPYCONTINUE parameter.

The combined total number of storage pools that are specified in the COPYSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the COPYCONTINUE value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to copy storage pools during the following operations:

- Back up and archive operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API
- Migration operations by IBM Spectrum Protect for Space Management clients
- Import operations that involve copying exported file data from external media to a storage pool defined with a copy storage pool list

Restrictions:

1. This parameter is available only to primary storage pools that use NATIVE or NONBLOCK data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Writing data simultaneously to copy storage pools is not supported when LAN-free data movement is used. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is accepted.
3. The simultaneous-write function is not supported for NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.
4. You cannot use the simultaneous-write function with CENTERA storage devices.

Attention: The function that is provided by the COPYSTGPOOLS parameter is not intended to replace the BACKUP STGPOOL command. If you use the COPYSTGPOOLS parameter, continue to use the BACKUP STGPOOL command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

COPYContinue

Specifies how the server reacts to a copy storage pool write failure for any of the copy storage pools that are listed in the COPYSTGPOOLS parameter. This parameter is optional. The default value is YES. When you specify the COPYCONTINUE parameter, you must also specify the COPYSTGPOOLS parameter.

The COPYCONTINUE parameter has no effect on the simultaneous-write function during migration.

You can specify the following values:

Yes

If the COPYCONTINUE parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

No

If the COPYCONTINUE parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the COPYCONTINUE parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The ACTIVEDATAPOOLS parameter is optional. Spaces between the names of the active-data pools are not allowed.

The combined total number of storage pools that are specified in the COPYSGTPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to active-data pools only during backup operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API.

Restrictions:

1. This parameter is available only to primary storage pools that use NATIVE or NONBLOCK data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Write data simultaneously to active-data pools is not supported when LAN-free data movement is used. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is accepted.
3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the TOCDESTINATION in the copy group of the management class has active-data pools defined, the active-data pools are ignored, and the data is stored into the primary storage pool only.
4. You cannot use the simultaneous-write function with CENTERA storage devices.
5. Data being imported is not stored in active-data pools. After an import operation, use the COPY ACTIVEDATA command to store the imported data in an active-data pool.

Attention: The function that is provided by the ACTIVEDATAPOOLS parameter is not intended to replace the COPY ACTIVEDATA command. If you use the ACTIVEDATAPOOLS parameter, use the COPY ACTIVEDATA command to ensure that the active-data pools contain all active data of the primary storage pool.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class. The default value is NO.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value 0 - 50. The default value is 1. If the value of the DEDuplicate parameter is NO, the default setting for IDENTIFYPROCESS has no effect.

Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Define a primary storage pool with an 8MMTAPE device class

Define a primary storage pool that is named 8MMPool to the 8MMTAPE device class (with a device type of 8MM) with a maximum file size of 5 MB. Store any files larger than 5 MB in subordinate pools, beginning with POOL1. Enable collocation of files for client nodes. Allow as many as 5 scratch volumes for this storage pool.

```
define stgpool 8mmpool 8mmtape maxsize=5m
  nextstgpool=pool1 collocate=node
  maxscratch=5
```

Related reference:

SET DRMDBBACKUPEXPIREDDAYS (Specify DB backup series expiration)

DEFINE STGPOOL (Define a copy storage pool assigned to sequential access devices)

Use this command to define a copy storage pool that is assigned to sequential access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--device_class_name----->
>>-POOLtype---Copy--+-----+----->
      '-DESCRIPTION---description-'
      .-ACCESS---READWrite-----
>--+-----+----->
      '-ACCESS---+READWrite---+'
          +-READOnly----+
          '-UNAVailable-'
      .-COLlocate---No----- .-RECLaim---100-----
>--+-----+-----+----->
      '-COLlocate---+No-----+' '-RECLaim---percent-'
          +-GRoup-----+
          +-NODE-----+
          '-FILESpace-'
      .-RECLAIMPRocess---1-----
>--+-----+----->
      '-RECLAIMPRocess---number-'
      .-RECLAMATIONType---THRESHold-----
>--+-----+----->
      |                                     (1) |
      '-RECLAMATIONType---+THRESHold+-----'
          '-SNAPlock--'
      .-OFFSITERECLAIMLimit---NOLimit-.
>--+-----+-----+----->
      '-OFFSITERECLAIMLimit---number--'
      .-REUsedelay---0-----
>--+-----+-----+----->
```


Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required.

You can specify one of the following options:

No

Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces named A, B, C, D, and E. File spaces A and B belong to a file space collocation group but C, D, and E do not. File spaces A and B are collocated by file space collocation group, while C, D, and E are collocated by node.

Data is collocated on the least amount of sequential access volumes.

NODe

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

FIlespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

REClaim

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the IBM Spectrum Protect database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100. The default value is 100, which means that reclamation is not completed.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

If you change the value from the default, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the unexpired files on the reclaimable volume from a primary or copy storage pool that is onsite. The process then writes these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location. However, the files can be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

RECLAIMPRocess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999. The default value is 1.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

You can specify one or more reclamation processes for each copy storage pool. You can specify multiple concurrent reclamation processes for a single copy storage pool, which makes better use of your available tape drives or FILE volumes. If multiple concurrent processing is not necessary, specify a value of 1 for the RECLAIMPROCESS parameter.

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. The following are possible values:

THRESHold

Specifies that volumes that belong to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes that belong to this storage pool are managed for retention by using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools that being defined to a server that has data retention protection that is enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, which is derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume by using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date expires, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools that are being defined must be the same when defined to the same device class name. The DEFINE command fails if the RECLAMATIONTYPE parameter specified is different from what is defined for storage pools that are already defined to the device class name.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes that space is reclaimed from during reclamation for this storage pool. This parameter is optional. The default value is NOLIMIT. You can specify the following values:

NOLimit

Specifies that you want to reclaim the space in all of your offsite volumes.

number

Specifies the number of offsite volumes to reclaim space from. You can specify an integer 0 - 99999. A value of zero means that none of the offsite volumes are reclaimed.

Tip:

To determine the value for the OFFSITERECLAIMLIMIT, use the statistical information in the message that is issued at the end of the offsite volume reclamation operation. The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose a copy storage pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 will be reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 will be reclaimed.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer 0 - 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the copy storage pool and the corresponding estimated capacity for the copy storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the copy storage pool until the access mode is changed. An administrator can then query the server for empty, offsite scratch volumes and return them to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to files in the copy storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters.

Enclose the location name in quotation marks if the location name contains any blank characters.

DATAFormat

Specifies the data format to use to back up files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. You can specify the following values:

NATIVE

Specifies the data format is the native IBM Spectrum Protect server format and includes block headers.

NONBLOCK

Specifies the data format is the native IBM Spectrum Protect server format and does not include block headers. The default minimum block size on a volume that is associated with a FILE device class is 256 KB, regardless how much data is written to the volume. For certain tasks, you can minimize wasted space on storage volumes by specifying the NONBLOCK data format. For example, you can specify the NONBLOCK data format for the following tasks:

- Using content-management products
- Using the DIRMC client option to store directory information
- Migrating very small files by using IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows

In most situations, however, the NATIVE format is preferred.

NETAPPDump

Specifies that the data is in a NetApp dump format. Do not specify this data format for file system images that are in a dump format and that were backed up from a NetApp file server by using NDMP. The server does not complete storage pool reclamation or AUDIT VOLUME for a storage pool with DATAFORMAT=NETAPPDUMP. You can use the MOVE DATA command to move NDMP-generated data out of a volume if the volume must be reused.

CELERRADump

Specifies that the data is in an EMC Celerra dump format. Do not specify this data format for file system images that are in a dump format and that were backed up from an EMC Celerra file server by using NDMP. The server does not complete storage pool reclamation or AUDIT VOLUME for a storage pool with DATAFORMAT=CELERRADUMP. You can use the MOVE DATA command to move NDMP-generated data out of a volume if the volume must be reused.

NDMPDump

Specifies that the data is in a NAS vendor-specific backup format. Do not specify this data format for file system images that are in a backup format and that were backed up from a NAS file server other than a NetApp or EMC Celerra file server. The server does not complete storage pool reclamation or AUDIT VOLUME for a storage pool with DATAFORMAT=NDMPDUMP. You can use the MOVE DATA command to move NDMP-generated data out of a volume if the volume must be reused.

CRCData

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCData to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class. The default value is NO.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value 0 - 50.

The default value for this parameter is 0. Data-deduplication processes for a copy storage pool are not necessary if you specify data-deduplication processes for the primary storage pool. When IBM Spectrum Protect analyzes a file in a storage pool, IBM Spectrum Protect also analyzes the file in all other storage pools.

Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Define a copy storage pool with a DC480 device class.

Define a copy storage pool, TAPEPOOL2, to the DC480 device class. Allow up to 50 scratch volumes for this pool. Delay the reuse of volumes for 45 days.

```
define stgpool tapepool2 dc480 pooltype=copy
maxscratch=50 reusedelay=45
```

Related reference:

SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)

DEFINE STGPOOL (Define an active-data pool assigned to sequential-access devices)

Use this command to define an active-data pool assigned to sequential-access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine STGpool--pool_name--device_class_name----->
>>-POOLtype----ACTIVEdata--+-+-----+----->
                                '-DESCRIPTION----description-'
    .-ACCESS----READWrite-----
>>-+-+-----+----->
    '-ACCESS----+READWrite----+'
        +-READOnly----+
        '-UNAVailable-'
    .-COLlocate----No----- .-REclaim----60-----
>>-+-+-----+----->
    '-COLlocate----+No-----+' '-REclaim----percent-'
        +-GRoup-----+
        +-NODE-----+
        '-Filespace-'
    .-RECLAIMProcess----1-----
>>-+-+-----+----->
    '-RECLAIMProcess----number-'
    .-RECLAMATIONType----THRESHold-----
>>-+-+-----+----->
    |                                     (1) |
```


COLlocate

Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required.

You can specify one of the following options:

No

Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces named A, B, C, D, and E. File spaces A and B belong to a file space collocation group but C, D, and E do not. File spaces A and B are collocated by file space collocation group, while C, D, and E are collocated by node.

Data is collocated on the least amount of sequential access volumes.

NODe

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

FIlespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

REClaim

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the IBM Spectrum Protect database.

Reclamation makes the fragmented space and space occupied by inactive backup files on volumes usable again by moving any remaining unexpired files and active backup files from one volume to another volume. This action makes the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100. The default value is 60.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

If you change the value from the default, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When an active-data pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the unexpired files on the reclaimable volume from a primary or active-data pool that is onsite. The process then writes these files to an available volume in the original active-data pool. Effectively, these files are moved back to the onsite location. However, the files can be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with active-data pools.

RECLAIMPRocess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999. The default value is 1.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

You can specify one or more reclamation processes for each active-data pool. You can specify multiple concurrent reclamation processes for a single active-data pool, which makes better use of your available tape drives or FILE volumes. If multiple concurrent processing is not necessary, specify a value of 1 for the RECLAIMPROCESS parameter.

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. The following are possible values:

THRESHold

Specifies that volumes that belong to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes that belong to this storage pool are managed for retention by using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools that are being defined to a server that has data retention protection that is enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, which is derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume by using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date expires, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools that are being defined must be the same when defined to the same device class name. The DEFINE command fails if the RECLAMATIONTYPE parameter specified is different from what is defined for storage pools that are already defined to the device class name.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes that space is reclaimed from during reclamation for this storage pool. This parameter is optional. The default value is NOLIMIT. You can specify the following values:

NOLimit

Specifies that you want to reclaim the space in all of your offsite volumes.

number

Specifies the number of offsite volumes to reclaim space from. You can specify an integer 0 - 99999. A value of zero means that none of the offsite volumes are reclaimed.

Tip:

To determine the value for the OFFSITERECLAIMLIMIT, use the statistical information in the message that is issued at the end of the offsite volume reclamation operation. The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose an active-data pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes are reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 are reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 is reclaimed.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer 0 - 10000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the active-data pool and the corresponding estimated capacity for the active-data pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the active-data pool until the access mode is changed. An administrator can then query the server for empty, offsite scratch volumes and return them to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to files in the active-data pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters.

Enclose the location name in quotation marks if the location name contains any blank characters.

DATAFormat

Specifies the data format to use to copy files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. You can specify the following values:

NATIVE

Specifies the data format is the native IBM Spectrum Protect server format and includes block headers.

NONblock

Specifies the data format is the native IBM Spectrum Protect server format and does not include block headers. The default minimum block size on a volume that is associated with a FILE device class is 256 KB, regardless how much data is written to the volume. For certain tasks, you can minimize wasted space on storage volumes by specifying the NONBLOCK data format. For example, you can specify the NONBLOCK data format for the following tasks:

- Using content-management products
- Using the DIRMC client option to store directory information
- Migrating very small files by using IBM Spectrum Protect for Space Management or IBM Spectrum Protect HSM for Windows

In most situations, however, the NATIVE format is preferred.

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. The default value is NO.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value 0 - 50.

The default value for this parameter is 0. Data-deduplication processes for a copy storage pool are not necessary if you specify data-deduplication processes for the primary storage pool. When IBM Spectrum Protect analyzes a file in a storage pool, IBM Spectrum Protect also analyzes the file in all other storage pools.

Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Define an active-data pool with a DC500 device class

Define an active-data pool, TAPEPOOL2, to the DC500 device class. Allow up to 50 scratch volumes for this pool. Delay the reuse of volumes for 45 days.

```
define stgpool tapepool3 dc500 pooltype=activedata
maxscratch=50 reusedelay=45
```

Related reference:

SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)

DEFINE STGPOOLDIRECTORY (Define a storage pool directory)

Use this command to define one or more directories in a directory-container or cloud-container storage pool.

Tip: After you define a cloud-container storage pool, create one or more directories that are used for local storage. You can temporarily store data in local storage during the data ingestion, before the data is moved to the cloud. In this way, you can improve backup and archive performance. For more information, see [Optimizing performance for cloud object storage](#).

Privilege class

To issue this command, you must have system privilege.

Syntax

```

          .-,-----
          v          |
>>-DEfine STGPOOLDIRectory--pool_name-----directory_name+-----><
```

Parameters

pool_name (Required)

Specifies the name of a directory-container or cloud-container storage pool. This parameter is required.

directory_name (Required)

Specifies the directory to be defined in the storage pool. This parameter is required. You can specify more than one directory name by separating each name with a comma, with no intervening spaces.

If you use the administrative client and the directory name contains a comma or a backslash ("\"), enclose the name in quotation marks.

Example: Define a storage pool directory

Define a storage pool directory that is named DIR1 by using a directory-container storage pool that is named POOL1.

AIX | Linux

```
define stgpooldirectory pool1 /storage/dir1
```

Windows

```
define stgpooldirectory pool1 c:\storage\dir1
```

Example: Define multiple storage pool directories

Define storage pool directories that are named DIR1 and DIR2 by using a directory-container storage pool that is named POOL1.

AIX | Linux

```
define stgpooldirectory pool1 /storage/dir1,/storage/dir2
```

Windows

```
define stgpooldirectory pool1 e:\storage\dir1,f:\storage\dir2
```

Example: Define local storage for a cloud-container storage pool

Create a storage pool directory that is named DIR3 in a cloud-container storage pool that is named CLOUDLOCALDISK1.

AIX Linux

```
define stgpooldirectory cloudlocaldisk1 /storage/dir3
```

Windows

```
define stgpooldirectory cloudlocaldisk1 c:\storage\dir3
```

Table 1. Commands related to DEFINE STGPOOLDIRECTORY

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE STGPOOLDIRECTORY	Deletes a storage pool directory from a directory-container or cloud-container storage pool.
QUERY STGPOOLDIRECTORY	Displays information about storage pool directories.
UPDATE STGPOOLDIRECTORY	Changes the attributes of a storage pool directory.

DEFINE STGRULE (Define a storage rule)

Use this command to define a storage rule.

The DEFINE STGRULE command takes several forms. The syntax and parameters for each form are defined separately.

Table 1. Commands related to DEFINE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (auditing)	Updates a storage rule for auditing storage pools.
UPDATE STGRULE (data deduplication statistics)	Updates a storage rule for generating data deduplication statistics.
UPDATE STGRULE (reclaiming)	Updates a storage rule for reclaiming cloud-container storage pools.
UPDATE STGRULE (tiering)	Updates a tiering storage rule.

- **DEFINE STGRULE (Define a rule for auditing storage pools)**
Use this command to schedule audit operations for a storage pool. The audit operations are designed to identify corrupted files within the storage pool.
- **DEFINE STGRULE (Define a rule for generating data deduplication statistics)**
Use this command to define a rule for generating data deduplication statistics. You can define one or more storage rules for a target container storage pool.
- **DEFINE STGRULE (Define a rule for reclaiming cloud containers)**
Use this command to define a rule for daily space reclamation in cloud-container storage pools. You can define one storage rule per storage pool.
- **DEFINE STGRULE (Define a storage rule for tiering)**
Use this command to define a storage rule for one or more storage pools. The storage rule schedules tiering between container storage pools. You can define one or more storage rules for a target container storage pool.

DEFINE STGRULE (Define a rule for auditing storage pools)

Use this command to schedule audit operations for a storage pool. The audit operations are designed to identify corrupted files within the storage pool.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-DEFine STGRULE--rule_name--storage_pool----->
                                     .-DELAY----7-----.
>----ACTiontype----AUDit-----+-----+----->
                                     '-DELAY----delay-'

    .-AUDITType----Extent-.    .-AUDITLevel----5-----.
>--+-----+-----+-----+-----+----->
                                     '-AUDITLevel----1+-'
                                     '-5-'

    .-STARTTime----current_time-.    .-ACTIVE----Yes-----.
>--+-----+-----+-----+-----+----->
    '-STARTTime----time-----'    '-ACTIVE----No----'
                                     '-Yes-'

>--+-----+-----+-----+-----+----->>
    '-DESCription----description-'

```

Parameters

rule_name (Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

storage_pool (Required)

Specifies the name of the storage pool to audit.

ACTiontype=AUDit (Required)

Specifies that the storage rule is for an audit operation.

DELAY

Specifies the interval, in days, between audit operations. This parameter is optional. The default value is 7 days. You can specify an integer in the range 1 - 9999.

AUDITType

Specifies the audit type. This parameter is optional. You can specify the following value:

Extent

Specifies that only extents are audited. This is the default value.

Restriction: In IBM Spectrum Protect™ Version 8.1.5, you can use the DEFINE STGRULE command with the ACTIONTYPE=AUDIT setting only to audit extents. Objects are not audited.

AUDITLevel

Specifies the level of the audit. This parameter is optional. The following values are possible:

1

Specifies a minimal audit operation of the extents in the storage pool.

5

Specifies a full audit operation of the extents in the storage pool. This is the default value.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. The default is the current time. This parameter is optional.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether storage rule processing occurs. This parameter is optional. The default is YES. The following values are possible:

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

DEscription

Specifies a description of the storage rule. This parameter is optional. The maximum length of the description is 255 characters. If the description includes spaces, enclose the description in quotation marks.

Define a rule for an extent-level audit operation

Define a storage rule, FULLAUDIT, to schedule a full audit of extents in storage pool DIRPOOL. The audit operation is started now and is repeated every three days:

```
define stgrule fullaudit dirpool actiontype=audit delay=3 auditlevel=5 starttime=now
```

Related commands

Table 1. Commands related to DEFINE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (auditing)	Updates a storage rule for auditing storage pools.

DEFINE STGRULE (Define a rule for generating data deduplication statistics)

Use this command to define a rule for generating data deduplication statistics. You can define one or more storage rules for a target container storage pool.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>--DEFine STGRULE--rule_name--target_stgpool----->
                                     .-DELAY---1-----
>----ACTiontype-----GENdedupstats-----+----->
                                     '-DELAY---delay-'
                                     .-MAXPProcess---8----- .-STARTTime---current_time-
>--+-----+-----+-----+----->
   '-MAXPProcess---number-' '-STARTTime---time-----'
                                     .-ACTIVE---Yes-----
>--+-----+-----+-----+----->
   '-ACTIVE---+No--+-'
                                     '-Yes-'
                                     .-NODEList---*-----
>--+-----+-----+-----+----->
   |                                     .-,-----|
   |                                     V         |
   '-NODEList---+node_name-----+--+-'
                                     '-node_group_name-'
                                     .-NAMEType---SERVER-----
```

```

>----->
'-NAMEType--+-SERVER--+'
      +-UNICODE+
      '-FSID----'

.-FSLIST-----*-----
>----->
|
|          v          |
|'-FSLIST--+-+filesystem_name-+-+'|
|          +-----+
|          '-fsid-----'|

.-CODEType--+-BOTH-----
>----->
'-CODEType--+-UNICODE-----+'
      +-NONUNICODE+
      '-BOTH-----'

>-----><
'-DESCRIPTION--+-description-'

```

Parameters

rule_name (Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

target_stgpool (Required)

Specifies the name of the target storage pool.

ACTiontype=GENdedupstats (Required)

Specifies that data deduplication statistics are generated.

DELAY

Specifies the interval, in days, between operations to collect statistics. The default value is 1 day. You can specify an integer in the range 0 - 9999.

MAXProcess

Specifies the maximum number of parallel processes to collect statistics. This parameter is optional. You can enter a value in the range 1 - 99. The default value is 8. For example, if you have 4 storage pools and you specify the default value for this parameter, 32 processes are started.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. The default is the current time. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether the storage rule processing occurs. This parameter is optional. The default is YES. The following values are possible:

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

NODEList

Specifies the name of the client node or defined group of client nodes for which data deduplication statistics are collected. You can also specify a combination of client node names and client-node group names. To specify multiple client node names or client-node group names, separate the names with commas with no intervening spaces. You can use wildcard

characters with client node names but not with client-node group names. The specified value can have a maximum of 1024 characters. The default value is an asterisk (*), which shows information for all client nodes.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. Use this parameter when IBM Spectrum Protect™ clients have file spaces that are in Unicode format, and are on Windows, NetWare, or Macintosh OS X operating systems. This parameter is optional.

This parameter is required if you specify a node name and a file space name or FSID.

Restriction: When you specify this parameter, the file space name cannot contain an asterisk.

Specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names. This is the default.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server's code page.

Tip: Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their FSIDs.

FSLIST

Specifies the names of one or more file spaces for which data deduplication statistics are collected. This parameter is optional. You can use wildcard characters to specify this name. The specified value can have a maximum of 1024 characters. An asterisk is the default. You can specify one of the following values:

*

Specify an asterisk (*) to show information for all file spaces or IDs.

filespace_name

Specifies the name of the file space. You can specify more than one file space by separating the names with commas and no intervening spaces.

fsid

Specifies the name of a file space identifier. This parameter is valid for clients with file spaces that are in Unicode format. Specify more than one file space by separating the names with commas and no intervening spaces.

For clients with file spaces that are in Unicode format, you can enter either a file space name or an FSID. If you enter a file space name, the server might have to convert the file space name that you enter. For example, the server might have to convert the name that you enter from the server's code page to Unicode.

Restrictions: The following restrictions apply to file space names and FSIDs:

- You must specify a node name if you specify a file space name.
- Do not specify both file space names and FSIDs on the same command.

CODEType

Specifies what type of file spaces to include in the record. The default value is BOTH, which specifies that file spaces are included regardless of code page type. Use this parameter only when you enter an asterisk to display information about all file spaces. This parameter is optional. Specify one of the following values:

UNICODE

Include file spaces that are in Unicode format.

NONUNICODE

Include file spaces that are not in Unicode format.

BOTH

Include file spaces regardless of code page type. This is the default.

DESCRIPTION

Specifies a description of the storage rule. This parameter is optional.

Define a rule to generate data deduplication statistics

Define a storage rule that is named MYSTAT1 to generate data deduplication statistics for the target storage pool, TARGET1. Limit the scope to a node that is named NODE1 and to the MYNODEGROUP node group. Limit the file spaces to FS1 and to all file spaces whose names start with FILESPACE1:


```
define stgrule mystat1 target1 actiontype=gendedupstats
nodelist=nodel,mynodegroup fslist=/fs1,/filespace1*
```

Related commands

Table 1. Commands related to DEFINE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (data deduplication statistics)	Updates a storage rule for generating data deduplication statistics.

DEFINE STGRULE (Define a rule for reclaiming cloud containers)

Use this command to define a rule for daily space reclamation in cloud-container storage pools. You can define one storage rule per storage pool.

Privilege class

To issue this command, you must have system privilege.

Restriction: You can configure a cloud reclamation rule for a storage pool only on a Microsoft Azure cloud computing system or on a cloud computing system with the Simple Storage Service (S3) protocol.

Syntax

```
>>-DEFine STGRULE--rule_name--pool_name----->
                                     .-PCTUnused---70-----.
>----ACTiontype----REClaim---+-----+----->
                                     '-PCTUnused---percentage-'
                                     .-MAXProcess---16----- .-DUration---120-----.
>--+-----+-----+-----+----->
   '-MAXProcess---number-' '-DUration---minutes-'
                                     .-STARTTime---current_time-. .-ACTIVE---Yes-----.
>--+-----+-----+-----+----->
   '-STARTTime---time-----' '-ACTIVE---+No---+'
                                     '-Yes-'
>--+-----+-----+-----+-----><
   '-DESCription---description-'
```

Parameters

rule_name (Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

pool_name (Required)

Specifies the name of the cloud-container storage pool.

ACTiontype=REClaim (Required)

Specifies that a cloud-container storage pool is reclaimed. Used data extents are moved to a new container. Unused extents are discarded.

PCTUnused

Specifies the percentage of the container that is no longer in use. After unused space reaches a percentage that you designate, the cloud container is reclaimed. The default value is 70 percent. You can specify an integer in the range 50 - 99. This parameter is optional.

MAXProcess

Specifies the maximum number of parallel processes that can be used to complete the storage rule for the storage pool that is specified. This parameter is optional. You can enter a value in the range 1 - 99. The default value is 16.

DUration

Specifies the maximum number of minutes that the storage rule runs before it is automatically canceled. You can specify a number in the range 60 - 1440. The default value is 120 minutes (2 hours). This parameter is optional.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. The default is the current time. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether the storage rule processing occurs. This parameter is optional. The default is YES. The following values are possible:

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

DEscription

Specifies a description of the storage rule. This parameter is optional.

Define a rule to reclaim space in a cloud-container storage pool

Define a storage rule that is named RECLAIMCTR1 to reclaim cloud containers that are more than half unused in storage pool CLOUDPOOL1. Specify a start time of 04:00 hours with a maximum of 2 processes for the storage rule:

```
define stgrule reclaimctr1 cloudpool1 actiontype=reclaim
pctunused=51 maxprocess=2 starttime=04:00:00
```

Related commands

Table 1. Commands related to DEFINE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (reclaiming)	Updates a storage rule for reclaiming cloud-container storage pools.

DEFINE STGRULE (Define a storage rule for tiering)


Use this command to define a storage rule for one or more storage pools. The storage rule schedules tiering between container storage pools. You can define one or more storage rules for a target container storage pool.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>--DEFine STGRULE--rule_name--target_stgpool----->
```



```

>----ACTiontype-----Tier-----SRCpools-----source_pool+----->
  .-TIERDelay-----30----- .-MAXPRocess-----8----- .
>--+-----+-----+-----+-----+-----+-----+----->
  '-TIERDelay-----delay-' '-MAXPRocess-----number-'

  .-DURation-----NOLimit-. .-STARTTime-----current_time-.
>--+-----+-----+-----+-----+-----+-----+----->
  '-DURation-----minutes-' '-STARTTime-----time-----'

  .-ACTIVE-----Yes----- .
>--+-----+-----+-----+-----+-----+-----+----->>
  '-ACTIVE-----+No--+-' '-DESCRiption-----description-'
    '-Yes-'

```

Parameters

rule_name(Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

target_stgpool(Required)

Specifies the name of the target cloud-container storage pool.

ACTiontype=Tier(Required)

Specifies that the storage rule tiers objects from the source storage pool to the target storage pool.

You can use tiering to lower storage costs by moving data to a cloud-container storage pool.

SRCPools(Required)

Specifies the name of the source directory-container storage pools. If you specify a pool as the source of a storage rule, you cannot specify the same pool as the source of another storage rule. To specify multiple storage pools, separate the names with commas with no intervening spaces. You must specify this parameter if the ACTIONTYPE=TIER parameter is specified.

TIERDelay

Specifies the number of days to wait before the storage rule tiers objects to the next storage pool. The default value is 30 days. You can specify an integer in the range 0 - 9999. The parameter value applies to all files in the storage pool.

MAXProcess

Specifies the maximum number of parallel processes to complete the storage rule for each source storage pool that is specified. This parameter is optional. Enter a value in the range 1 - 99. The default value is 8. For example, if you have 4 source storage pools and you specify the default value for this parameter, 32 processes are started.

DURation

Specifies the maximum number of minutes that the storage rule runs before it is automatically canceled. You can specify a number in the range 60 - 1440. The default value is unlimited. If you do not specify a value, or if you specify a value of NOLimit, the storage rule runs until it is completed. This parameter is optional.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. The default is the current time. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

Specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether the storage rule processing occurs. This parameter is optional. The default is YES. The following values are possible:

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

DESCRiption

Specifies a description of the storage rule. This parameter is optional.

Define a storage rule

Define a storage rule that is named tieraction to move data from the source directory-container storage pools dirpool1 and dirpool2 to the target cloud-container storage pool cloudpool1. Specify a start time of 03:00 hours that uses a maximum of 10 processes for a tiering storage rule:

```
define stgrule tieraction cloudpool1 srcpools=dirpool1,dirpool2
actiontype=tier maxprocess=10 starttime=03:00:00
```

Related commands

Table 1. Commands related to DEFINE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (tiering)	Updates a tiering storage rule.

DEFINE SUBSCRIPTION (Define a profile subscription)

Use this command on a managed server to subscribe that managed server to a profile.

When a server subscribes to its first profile, a subscription is also created to the default profile (if one exists) of the configuration manager. The server then contacts the configuration manager periodically for configuration updates.

Restrictions:

1. A server cannot subscribe to profiles from more than one configuration manager.
2. If a server subscribes to a profile with an associated object that is already defined on the server, the local definition is replaced by the definition from the configuration manager. For example, if a server has an administrative schedule named WEEKLY_BACKUP, then subscribes to a profile that also has an administrative schedule named WEEKLY_BACKUP, the local definition is replaced.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DEFine SUBSCRIPtion--profile_name----->
>--+-----+----->>
  '-SERVer----server_name-'
```

Parameters

profile_name (Required)

Specifies the name of the profile to which the server subscribes.

SERVer

Specifies the name of the configuration manager from which the configuration information is obtained. This parameter is required, if the managed server does not have at least one subscription. If the managed server has a subscription, you can omit this parameter and it defaults to the configuration manager for that subscription.

Example: Define a profile subscription

Subscribe a profile named BETA that resides on a configuration manager named TOM.

```
define subscription beta server=tom
```

Related commands

Table 1. Commands related to DEFINE SUBSCRIPTION

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFILE	Deletes a profile from a configuration manager.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGREFRESH	Specifies a time interval for managed servers to contact configuration managers.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)

Use this command to define a virtual file space mapping.

Virtual file space names can be used in the NAS data operations BACKUP NODE and RESTORE NODE similar to a file system name. Refer to the documentation about your NAS device for guidance on specifying the parameters for this command.

Note: The NAS node must have an associated data mover definition because when the IBM Spectrum Protect™ server updates a virtual file space mapping, the server attempts to contact the NAS device to validate the virtual file system and file system name.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned.

Syntax

```
>>-DEFine VIRTUALFSmapping -node_name----->
>>--virtual_filespace_name--file_system_name--path----->
  .-NAMEType----SERVER-----
>--+-----+-----+----->>
  '-NAMEType----+SERVER-----'
    '-HEXadecimal-'
```

Parameters

node_name (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

virtual_filespace_name (Required)

Specifies the name which refers to this virtual file space definition. The virtual file space name is case sensitive and the first character must be a forward slash /. The length of the name cannot be more than 64 characters, including the required forward slash. Virtual file space names are restricted to the same character set as all other objects in the server except that the forward slash / character is also allowed.

The virtual file space name cannot be identical to any file system on the NAS node. When selecting a virtual file space name, consider the following restrictions:

- If a file system is created on the NAS device with the same name as a virtual file system, a name conflict will occur on the server when the new file space is backed up. Use a string for the virtual file space name that is unlikely to be used as a real file system name on your NAS device in the future.

For example: A user follows a naming convention for creating file spaces on a NAS device with names of the form /vol1, /vol2, /vol3. The user defines a virtual file space to the server with the name /vol9. If the user continues to use the same naming convention, the virtual file space name is likely to conflict with a real file space name at some point in the future.

- During backup and restore operations, the server verifies that a name conflict does not occur prior to starting the operation.
- The virtual file space name appears as a file space in the output of the QUERY FILESPACE command, and also in the backup and restore panels of the IBM Spectrum Protect web client. Therefore, consider selecting a name that unambiguously identifies this object as a directory path on the NAS device.

file_system_name (Required)

Specifies the name of the file system in which the path is located. The file system name must exist on the specified NAS node. The file system name cannot contain wildcard characters.

path (Required)

Specifies the path from the root of the file system to the directory. The path can only reference a directory. The maximum length of the path is 1024 characters. The path name is case sensitive.

NAMEType

Specifies how the server should interpret the path name specified. This parameter is useful when a path contains characters that are not part of the code page in which the server is running. The default value is SERVER.

Possible values are:

SERVER

The server uses the server code page to interpret the path name.

HEXadecimal

The server interprets the path that you enter as the hexadecimal representation of the path. This option should be used when a path contains characters that cannot be entered. This could occur if the NAS file system is set to a language different from the one in which the server is running.

Example: Define a virtual file space mapping

Define the virtual file space mapping name /mikeshomedir for the path /home/mike on the file system /vol/vol1 on the NAS node named NAS1.

```
define virtualfsmapping nas1 /mikeshomedir /vol/vol1 /home/mike
```

Related commands

Table 1. Commands related to DEFINE VIRTUALFSMAPPING

Command	Description
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
QUERY VIRTUALFSMAPPING	Query a virtual file space mapping.

Command	Description
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

DEFINE VOLUME (Define a volume in a storage pool)

Use this command to assign a random or sequential access volume to a storage pool.

When you define a random-access (DISK) storage-pool volume or a sequential access storage pool volume that is associated with a FILE device class, you can have the server create the volume before it is assigned. Alternatively, you can use space triggers to create preassigned volumes when predetermined space-utilization thresholds are exceeded. For details about space triggers, see DEFINE SPACETRIGGER (Define the space trigger). For volumes associated with device classes other than DISK or device types other than FILE, you can use the DEFINE VOLUME command to assign an already-created volume to a storage pool.

AIX **Linux** When you use a FILE device class for storage that is managed by a z/OS® media server, it is not necessary to format or define volumes. If you define a volume for such a FILE device class by using the DEFINE VOLUME command, the z/OS media server does not allocate space for the volume until the volume is opened for its first use.

Attention: Volumes for the z/OS media server that are created using the DEFINE VOLUME command remain physically full or allocated after the server empties the volume, for example, after expiration or reclamation. For FILE volumes, the DASD space is not relinquished to the system when the volume is emptied. If a storage pool requires an empty or filling volume, the FILE volume can be used. In contrast, tape volumes that are logically empty are the same as physically empty. FILE and tape volumes remain defined in the server. In contrast, SCRATCH volumes, including the physical storage that is allocated for SCRATCH FILE volumes, are returned to the system when emptied.

To create space in sequential access storage pools, you can define volumes or allow the server to request scratch volumes as needed, as specified by the MAXSCRATCH parameter for the storage pool. For storage pools associated with the FILE device class, the server can create private volumes as needed using storage-pool space triggers. For DISK storage pools, the scratch mechanism is not available. However, you can create space by creating volumes and then defining them to the server. Alternatively, you can have the server create volumes that use storage-pool space triggers.

The server does not validate the existence of a volume name when defining a volume in a storage pool that is associated with a library. The defined volume has "0" EST capacity until data is written to the volume.

Attention: The size of a storage pool volume cannot be changed after it is defined to the server.

AIX If you change the size of IBM Spectrum Protect™ volumes by extending raw logical volumes through SMIT or otherwise altering the file sizes of the volumes with operating system commands or utilities, the server might not initialize correctly and data can be lost.

Windows If you change the size of volumes by altering the file sizes of the volumes with operating system commands or utilities, the server might not initialize correctly and data can be lost.

Restrictions:

- You cannot use this command to define volumes in storage pools with the parameter setting RECLAMATIONTYPE=SNAPLOCK. Volumes in this type of storage pool are allocated by using the MAXSCRATCH parameter on the storage pool definition.
- You cannot define volumes in a storage pool that is defined with the CENTERA device class.
- Linux** You cannot use raw logical volumes for storage pool volumes.

Physical files that are allocated with DEFINE VOLUME command are not removed from a file space if you issue the DELETE VOLUME command.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is assigned.

Syntax

```
>>-DEFine Volume--pool_name--volume_name----->
```


Volume Name Requirements	Example
<p>The name of the file to contain the volume data, with either the fully qualified path name or a path name relative to the current working directory.</p> <p>Windows If a name contains embedded blanks, equal signs, or other special characters, enclose the list in quotation marks.</p>	<p>AIX Linux</p> <pre>/usr/storage/sbkup01.dsm</pre> <p>AIX If you are using an AIX® logical volume, enter the path name as:</p> <pre>/dev/rxxx</pre> <p>where xxx is the logical volume name.</p> <p>Windows</p> <pre>"c:\program files\tivoli\tsm\server\data3.dsm"</pre>

Table 2. Volume name requirements for FILE

Volume Name Requirements	Example
<p>The name of the file to contain the volume data, with either the fully qualified path name or the path name relative to a directory identified in the DIRECTORY parameter for the device class.</p> <p>Windows If a name contains embedded blanks, equal signs, or other special characters, enclose the list in quotation marks.</p> <p>Place FILE volumes in one of the directories that are specified with the DIRECTORY parameter of the DEFINE DEVCLASS command. Otherwise, storage agents might not have access to the volumes. For details, see DEFINE PATH (Define a path).</p>	<p>AIX Linux</p> <pre>/data/fpool01.dsm</pre> <p>Windows</p> <pre>"f:\data storage\fpool01.dsm"</pre>

Table 3. z/OS media server: Volume name requirements for FILE

Volume Name Requirements	Example
--------------------------	---------

Volume Name Requirements	Example
<p>For FILE volumes used with the z/OS media server server, specify a data set name. The data set name can consist of one or more qualifiers that are delimited by a period. The qualifiers can contain up to 8 characters. The maximum length of the data set name is 44 characters. The first letter of each qualifier must be alphabetic or national (@#\$), followed by alphabetic, national, hyphen, or numeric characters.</p> <p>To allocate the associated VSAM Linear Dataset when the volume is tendered on the z/OS system, the High Level Qualifier (HLQ) is typically filtered by specific ACS routines within the SMS policy constraints on the system where the z/OS media server is running.</p> <p>The behavior of the HLQ is similar to the behavior of the PREFIX name on a scratch request. The HLQ is typically used by DFSMS to affect allocation attributes, such as Extended Addressability for data sets that are expected to extend when space that is already allocated to the file volume is used up.</p> <p>If the data set does not exist, the server creates it when the volume is used for a specific IBM Spectrum Protect storage operation. The data set is not created when the volume is defined. Data loss can result when defining volumes because the z/OS media server reuses the volume or VSAM LDS if it exists at the time of allocation time.</p> <p>Important: To allow the server to generate volume names, consider using SCRATCH volumes.</p>	<div style="display: flex; justify-content: space-between; border-bottom: 1px solid black; padding-bottom: 2px;"> AIX Linux </div> <p>SERVER1.BFS.POOL3.VOLA</p>

Table 4. Volume name requirements for tape

Volume Name Requirements	Example
<p>Use 1 - 32 alphanumeric characters.</p> <p>The volume name cannot contain any embedded blanks or equal signs.</p>	<p>DSMT01</p>

AIX

Linux

Table 5. z/OS media server: Volume name requirements for tape

Volume Name Requirements	Example
<p>For tape cartridges, specify a tape volume name with 1 - 6 alphanumeric characters. The server converts tape volume names to uppercase.</p> <p>The volume name cannot contain any embedded blanks or equal signs.</p> <p>Each volume that is used by a server for any purpose must have a unique name. This requirement applies to all volumes, whether the volumes are used for storage pools, or used for operations such as database backup or export. The requirement also applies to volumes that reside in different z/OS media libraries but that are used by the same server.</p>	<p>DSMT01</p>

Table 6. Volume name requirements for REMOVABLEFILE

Volume Name Requirements	Example
--------------------------	---------

Volume Name Requirements	Example
1–6 alphanumeric characters	DSM01
The server converts volume names to uppercase.	

ACcESS

Specifies how client nodes and server processes (such as migration) can access files in the storage pool volume. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that client nodes and server processes can read from and write to files stored on the volume.

READOnly

Specifies that client nodes and server processes can only read files that are stored on the volume.

UNAVailable

Specifies that client nodes or server processes cannot access files that are stored on the volume.

If you define a random access volume as UNAVAILABLE, you cannot vary the volume online.

If you define a sequential access volume as UNAVAILABLE, the server does not attempt to access the volume.

OFFsite

Specifies that the volume is at an offsite location from which it cannot be mounted. You can specify this value only for volumes in copy or active-data storage pools.

Use this value to help you track volumes at offsite locations. The server treats volumes that are designated as offsite differently:

- The server does not generate mount requests for volumes designated offsite.
- The server reclaims or moves data from offsite volumes by retrieving files from other storage pools.
- The server does not automatically delete empty, offsite scratch volumes from a copy or active-data storage pool.

LOCation

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential access storage pools. The location information can be a maximum length of 255 characters. Enclose the location in quotation marks if it contains any blank characters.

FORMatsize

Specifies the size of the random access volume or FILE volume that is created and formatted in one step. The value is specified in megabytes. The maximum size is 8 000 000 MB (8 terabytes). This parameter is required if any of the following conditions are true:

- A single FILE or DISK volume is specified, which is to be created and formatted in one step.
- The value for the NUMBEROFVOLUMES parameter is greater than 1, and DISK volumes are being created.
- The value of the NUMBEROFVOLUMES parameter is greater than 1, and the value of the FORMATSIZE parameter is less than or equal to the MAXCAPACITY parameter of the DEFINE DEVCLASS command.

If you are allocating volumes on a z/OS media server, this parameter is not valid.

For a FILE volume, you must specify a value less than or equal to the value of the MAXCAPACITY parameter of the device class associated with the storage pool.

You cannot use this parameter for multiple, predefined volumes. Unless you specify `WAIT=YES` is specified, the operation is completed as a background process.

NUMBERofvolumes

Specifies the number of volumes that are created and formatted in one step. This parameter applies only to storage pools with DISK or FILE device classes. This parameter is optional. The default is 1. If you specify a value greater than 1, you must also specify a value for the FORMATSIZE parameter. Specify a number from 1 to 256.

If you are allocating volumes on a z/OS media server, the only value that this parameter supports is the default value of 1.

If the value for the NUMBEROFVOLUMES parameter is greater than 1, the volume name you specified will have a numeric suffix appended to create each name, for example, `tivolivol001` and `tivolivol002`. Be sure to choose a volume name so that a valid file name for the target file system is created when the suffix is appended.

Important: You must ensure that storage agents can access newly created FILE volumes. For more information, see `DEFINE PATH` (Define a path).

Wait

Specifies whether volume creation and formatting operation is completed in the foreground or background. This parameter is optional. It is ignored unless you also specify the FORMATSIZE parameter.

No

Specifies that a volume creation and formatting operation is completed in the background. The NO value is the default when you also specify a format size.

Yes

Specifies that a volume creation and formatting operation is completed in the foreground.
Remember: You cannot specify `WAIT=YES` from the server console.

Example: Use a background process to define a new 100 MB volume for a disk storage pool

Create a volume of 100 MB in the disk storage pool named BACKUPPOOL. AIX Linux The volume name is `/var/storage/bf.dsm`. Windows The volume name is `j:\storage\bf.dsm`. Let the volume be created as a background process.

AIX Linux

```
define volume backuppool  
/var/storage/bf.dsm formatsize=100
```

Windows

```
define volume backuppool j:\storage\bf.dsm formatsize=100
```

Example: Define a volume to a disk storage pool with read and write access

A storage pool named POOL1 is assigned to a tape device class. Define a volume named TAPE01 to this storage pool, with READWRITE access.

```
define volume pool1 tape01 access=readwrite
```

Example: Define a volume to a file storage pool

A storage pool that is named FILEPOOL is assigned to a device class with a device type of FILE. AIX Linux Define a volume that is named `filepool_vol01` to this storage pool. Windows Define a volume that is named `fp_vol01.dsm` to this storage pool. AIX Linux

```
define volume filepool /usr/storage/filepool_vol01
```

Windows

```
define volume filepool j:\storage\fp_vol01.dsm
```

Example: Example: Use a background process to define 10 volumes for a file storage pool with a device class 5 GB maximum capacity

Define 10 volumes in a sequential storage pool that uses a FILE device class. The storage pool is named FILEPOOL. The value of the MAXCAPACITY parameter for the device class that is associated with this storage pool is 5 GB. Creation must occur in the background.

```
define volume filepool filevol numberofvolumes=10 formatsize=5000
```

The server creates volume names `filevol001` through `filevol010`.

Volumes are created in the directory or directories that are specified with the DIRECTORY parameter of the device class that is associated with storage pool `filepool`. If you specified multiple directories for the device class, individual volumes can be created in any of the directories in the list.

Related commands

Table 7. Commands related to DEFINE VOLUME

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.

Command	Description
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

DELETE commands

Use the DELETE commands to delete or remove an IBM Spectrum Protect™ object.

- DELETE ASSOCIATION (Delete the node association to a schedule)
- DELETE ALERTTRIGGER (Remove a message from an alert trigger)
- DELETE BACKUPSET (Delete a backup set)
- DELETE CLIENTOPT (Delete an option in an option set)
- DELETE CLOPTSET (Delete a client option set)
- DELETE COLLOGROUP (Delete a collocation group)
- DELETE COLLOCMEMBER (Delete collocation group member)
- DELETE COPYGROUP (Delete a backup or archive copy group)
- DELETE DATAMOVER (Delete a data mover)
- DELETE DEDUPSTATS (Delete data deduplication statistics)
- DELETE DEVCLASS (Delete a device class)
- DELETE DOMAIN (Delete a policy domain)
- DELETE DRIVE (Delete a drive from a library)
- DELETE EVENT (Delete event records)
- DELETE EVENTSERVER (Delete the definition of the event server)
- DELETE FILESPACE (Delete client node data from the server)
- DELETE GRPMEMBER (Delete a server from a server group)
- DELETE LIBRARY (Delete a library)
- DELETE MACHINE (Delete machine information)
- DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)
- DELETE MGMTCLASS (Delete a management class)
- DELETE NODEGROUP (Delete a node group)
- DELETE NODEGROUPMEMBER (Delete node group member)
- DELETE PATH (Delete a path)
- DELETE POLICYSET (Delete a policy set)
- DELETE PROFASSOCIATION (Delete a profile association)
- DELETE PROFILE (Delete a profile)
- DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)
- DELETE RECOVERYMEDIA (Delete recovery media)
- DELETE SCHEDULE (Delete a client or an administrative command schedule)
- DELETE SCRIPT (Delete command lines from a script or delete the entire script)
- DELETE SERVER (Delete a server definition)
- DELETE SERVERGROUP (Delete a server group)
- DELETE SPACETRIGGER (Delete the storage pool space triggers)
- DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)
- DELETE STGRULE (Delete storage rules for storage pools)
- DELETE STGPOOL (Delete a storage pool)
- DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)
- DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)
- DELETE SUBSCRIPTION (Delete a profile subscription)
- DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)
- DELETE VOLHISTORY (Delete sequential volume history information)
- DELETE VOLUME (Delete a storage pool volume)

DELETE ALERTTRIGGER (Remove a message from an alert trigger)

Use this command to remove a message from the list of alert triggers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
      .-,-----  
      v |  
>>-DELeTe ALERtTrigger-----+---message_number+-----><
```

Parameters

message_number (Required)

Specifies the message number that you want to remove from the list of alert triggers. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers.

Delete alert trigger

Delete two message numbers that are designated as alerts, by issuing the following command:

```
delete alerttrigger ANR1067E,ANR1073E
```

Related commands

Table 1. Commands related to DELETE ALERTTRIGGER

Command	Description
DEFINE ALERTTRIGGER (Define an alert trigger)	Associates specified messages to an alert trigger.
QUERY ALERTSTATUS (Query the status of an alert)	Displays information about alerts that have been issued on the server.
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	Displays message numbers that trigger an alert.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
UPDATE ALERTTRIGGER (Update a defined alert trigger)	Updates the attributes of one or more alert triggers.
UPDATE ALERTSTATUS (Update the status of an alert)	Updates the status of a reported alert.

DELETE ASSOCIATION (Delete the node association to a schedule)

Use this command to delete the association of a client node to a client schedule. IBM Spectrum Protect™ no longer runs the schedule on the client node.

If you try to disassociate a client from a schedule to which it is not associated, this command has no effect for that client.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the schedule belongs

Syntax

```
>>-DELeTe ASSOCIation--domain_name--schedule_name----->  
  
      .-,-----
```

```
      v      |
>-----node_name+-----<<
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedule belongs.

schedule_name (Required)

Specifies the name of the schedule from which clients are to be disassociated.

node_name (Required)

Specifies the name of the client node that is no longer associated with the client schedule. You can specify a list of clients which are to be no longer associated with the specified schedule. Commas, with no intervening spaces, separate the items in the list. You can also use a wildcard character to specify a name. All matching clients are disassociated from the specified schedule.

Example: Delete a node association to a schedule

To delete the association of the node JEFF, assigned to the DOMAIN1 policy domain, to the WEEKLY_BACKUP schedule issue the following command:

```
delete association domain1 weekly_backup jeff
```

Example: Delete a node association to a schedule using a wildcard for node selection

Delete the association of selected clients, assigned to the DOMAIN1 policy domain, to the WEEKLY_BACKUP schedule so that this schedule is no longer run by these clients. The nodes that are disassociated from the schedule contain ABC or XYZ in the node name. Issue the command:

```
delete association domain1 weekly_backup *abc*,*xyz*
```

Related commands

Table 1. Commands related to DELETE ASSOCIATION

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.

DELETE BACKUPSET (Delete a backup set)

Use this command to manually delete a backup set before its retention period expires.

When the server creates a backup set, the retention period assigned to the backup set determines how long the backup set remains in the database. When that date passes, the server automatically deletes the backup set when expiration processing runs. However, you can also manually delete the client's backup set from the server before it is scheduled to expire by using the DELETE BACKUPSET command.

Attention: If the volumes contain multiple backup sets, they are not returned to scratch status until all the backup sets are expired or are deleted.

Privilege class

If the REQSYSAUTHOUTFILE server option is set to YES (the default), the administrator must have system privilege. If the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```
      v      |
      .-----.
```

```

>>-DElete BACKUPSET-----+node_name-----+----->
      '-node_group_name-'
      .-,------.
      v          |
>---backup_set_name+-----+----->
      '-BEGINDate----date-'
>-----+-----+-----+----->
      '-BEGINTime----time-' '-ENDDate----date-'
      .-WHEREDATAType----ALL-----
>-----+-----+-----+----->
      '-ENDTime----time-' | .-,------. |
      |                   v          | |
      '-WHEREDATAType-----+FILE--+--+'
      '-IMAGE-'
>-----+-----+-----+----->
      '-WHERERETention----+days--+-'
      '-NOLimit-'
>-----+-----+-----+----->
      '-WHEREDEScRiption----+description-'
      .-Preview ----No-----
>-----+-----+-----+----->>
      '-Preview----+No--+-'
      '-Yes-'

```

Parameters

node_name or node_group_name (Required)

Specifies the name of the client nodes or node groups whose data is contained in the specified backup set volumes. To specify multiple node and node group names, separate the names with commas and no intervening spaces. Any node name you specify may contain wildcard characters, but node group names cannot contain wildcard characters. If backup set volumes contain backup sets from multiple nodes then every backup set whose node name matches one of the specified node names will be deleted.

backup_set_name (Required)

Specifies the name of the backup set to delete. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

BEGINDate

Specifies the beginning date in which the backup set to delete was created. This parameter is optional. You can use this parameter with the BEGINTIME parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the BEGINDATE parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes specified	NOW+02:00 <i>or</i> +02:00.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes specified	NOW-02:00 <i>or</i> -02:00.

ENDDate

Specifies the ending date in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the ENDTIME parameter to specify a range for the date and time. If you specify an end date without an end time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days <i>or</i> +days	The current date plus days specified.	TODAY +3 <i>or</i> +3.
TODAY-days <i>or</i> -days	The current date minus days specified.	TODAY -3 <i>or</i> -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the ENDDATE parameter to specify a range for the date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified end date	NOW+02:00 <i>or</i> +02:00.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified end date	NOW-02:00 <i>or</i> -02:00.

WHERE DATATYPE

Specifies the backup sets containing the specified types of data are to be deleted. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be deleted. To specify multiple data types, separate the data types with commas and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be deleted. This is the default.

FILE

Specifies that a file level backup set is to be deleted. File level backup sets contain files and directories backup up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be deleted. Image backup sets contain images created by the backup-archive client BACKUP IMAGE command.

WHERERETention

Specifies the retention value, specified in days, that is associated with the backup sets to delete. You can specify an integer from 0 to 30000. The values are:

days

Specifies that backup sets that are retained this number of days are deleted.

NOLimit

Specifies that the backup sets that are retained indefinitely are deleted.

WHEREDESCRIPTION

Specifies the description that is associated with the backup set to delete. The description you specify can contain a wildcard character. This parameter is optional. Enclose the description in quotation marks if it contains any blank characters.

Preview

Specifies whether to preview the list of backup sets to delete, without actually deleting the backup sets. This parameter is optional. The default value is NO. The values are:

No

Specifies that the backup sets are deleted.

Yes

Specifies that the server displays the list of backup sets to delete, without actually deleting the backup sets.

Example: Delete a backup set

Delete backup set named PERS_DATA.3099 that belongs to client node JANE. The backup set was generated on 11/19/1998 at 10:30:05 and the description is "Documentation Shop".

```
delete backupset pers_data.3099
begindate=11/19/1998 begintime=10:30:05
wheredescription="documentation shop"
```

Related commands

Table 1. Commands related to DELETE BACKUPSET

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.

Command	Description
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DELETE CLIENTOPT (Delete an option in an option set)

Use this command to delete a client option in an option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege.

Syntax

```
>>-DELEte CLIENTOpt--option_set_name--option_name----->
>--+-----+----->>
  '-SEQnumber-----+number-+-'
                    '-ALL----'
```

Parameters

- option_set_name (Required)
Specifies the name of the client option set.
- option_name (Required)
Specifies a valid client option.
- SEQnumber
Specifies a sequence number when an option name is specified more than once. This parameter is optional. Valid values are:
 - n
Specifies an integer of 0 or greater.
 - ALL
Specifies all sequence numbers.

Example: Delete the date format option

Delete the date format option in an option set named *ENG*.

```
delete clientopt eng dateformat
```

Related commands

Table 1. Commands related to DELETE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DELETE CLOPTSET (Delete a client option set)

Use this command to delete a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege.

Syntax

```
>>-DELeTe CLOptset--option_set_name-----<<
```

Parameters

option_set_name (Required)
Specifies the name of the client option set to delete.

Example: Delete a client option set

Delete the client option set named ENG.

```
delete cloptset eng
```

Related commands

Table 1. Commands related to DELETE CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DELETE COLLOGROUP (Delete a collocation group)

Use this command to delete a collocation group. You cannot delete a collocation group if it has any members in it.

You can remove all the members in the collocation group by issuing the DELETE COLLOCMEMBER command with a wildcard in the node_name parameter.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
>>-DELeTe COLLOGroup--group_name-----<<
```

Parameters

group_name
Specifies the name of the collocation group that you want to delete.

Example: Delete a collocation group

Delete a collocation group named group1.

```
delete collogroup group1
```

Related commands

Table 1. Commands related to DELETE COLLOGROUP

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DELETE COLLOCMEMBER (Delete collocation group member)

Use this command to delete a client node or file space from a collocation group.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

Delete a node from a collocation group

```
                .-,-,-----  
                v          |  
>>-DELeTe COLLOCMemBer--group_name----node_name-+-----><
```

Parameters

group_name
Specifies the name of the collocation group from which you want to delete a client node.
node_name

Specifies the name of the client node that you want to delete from the collocation group. You can specify one or more names. When you specify multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple nodes.

Delete a file space from a file space collocation group

```
>>-DELEte COLLOCMember--group_name--node_name----->
      .-,------.
      v          |
>--Filespace--====file_space_name+----->
      .-NAMEType--==--SERVER-----
>--+-----+-----+----->
      '-NAMEType--==--+SERVER--+-'
          +-UNICODE-+
          '-FSID----'

      .-CODEType--==--BOTH-----
>--+-----+-----+-----><
      '-CODEType--==--+BOTH-----+'
          +-UNICODE-----+
          '-NONUNICODE-'
```

Parameters

group_name

Specifies the name of the collocation group from which you want to delete a file space.

node_name

Specifies the client node where the file space is located.

Filespace

Specifies the *file_space_name* on the client node that you want to delete from the collocation group. You can specify one or more file space names that are on a specific client node. If you specify multiple file space names, separate the names with commas, and do not use intervening spaces. You can also use wildcard characters when you specify multiple file space names.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with Unicode support. A backup-archive client with Unicode support is available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare. Use this parameter when you specify a file space name that is not a single wildcard. You can specify a fully qualified file space name, which does not have a wildcard. Or you can specify a partly qualified file space name, which can have a wildcard but must contain other characters. The default value is SERVER. Possible values are

SERVER

The server uses the server code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the names and the server code page. Conversion might fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names by their file space IDs (FSIDs).

CODEType

Specify how you want the server to interpret the file space names that you enter. Use this parameter only when you use a single wildcard character for the file space name. The default is BOTH, so the file spaces are included, regardless of code page type. The following values are available:

BOTH

Include the file spaces, regardless of code page type.

UNICODE

Include file spaces that are in Unicode only.

NONUNICODE
Include file spaces that are not in Unicode.

Delete collocation group members

Delete two nodes, NODE1 and NODE2, from a collocation group, GROUP1.

```
delete collocmember group1 node1,node2
```

Delete a file space from a file space collocation group

Issue the following command to delete file space *cap_27400* from collocation group *collgrp_2* on node *hp_4483*:

```
delete collocmember collgrp_2 hp_4483 filespace=cap_27400
```

Delete a file space collocation group member from a node that uses Unicode

If the file space is on a node that uses Unicode, you can specify that in the command. Issue the following command to delete file space *cap_257* from collocation group *collgrp_3* from the *win_4687* node:

```
delete collocmember collgrp_3 win_4687 filespace=cap_257 codetype=unicode
```

Delete a file space with a partial name designated

If the file space has a partial name, you can use a wildcard to delete it. Issue the following command to delete file space *cap_* from collocation group *collgrp_4* from *win_4687* node:

```
delete collocmember collgrp_4 win_4687 filespace=cap_* codetype=unicode
```

If there is more than one file space whose name begins with *cap_*, those file spaces are also deleted.

Related commands

Table 1. Commands related to DELETE COLLOCMEMBER

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

Table 1. Commands related to DELETE COPYGROUP

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
QUERY COPYGROUP	Displays the attributes of a copy group.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

DELETE DATAMOVER (Delete a data mover)

Use this command to delete a data mover. You cannot delete the data mover if any paths are defined for this data mover.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>--DELeTe DATAMover--data_mover_name-----><
```

Parameters

data_mover_name (Required)

Specifies the name of the data mover.

Note: This command deletes the data mover even if there is data for the corresponding NAS node.

Example: Delete a data mover

Delete the data mover for the node named NAS1.

```
delete datamover nas1
```

Related commands

Table 1. Commands related to DELETE DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE PATH	Defines a path from a source to a destination.
DELETE PATH	Deletes a path from a source to a destination.
QUERY DATAMOVER	Displays data mover definitions.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DATAMOVER	Changes the definition for a data mover.

AIX | Linux | Windows

DELETE DEDUPSTATS (Delete data deduplication statistics)

Use this command to delete data deduplication statistics for a directory-container storage pool or a cloud storage pool. You cannot delete the most recent data deduplication statistics for a client node and a file space.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool.

Syntax

```
>>-DELEte DEDUPStats--pool_name----->
                                     '-node_name-'

. -*----- . .-CODEType---BOTH-----
>-----+-----+-----+----->
| .-,----- . | '-CODEType---+UNICODE---'
| V             | |           +-NONUNICODE+
+---file_space_name---+         '-BOTH-----'
| .-,----- . |
| V             | |
'-----FSID-----'

.-NAMEType---SERVER-----
>-----+-----+-----+----->
'-NAMEType---+SERVER---' '-TODate---date-'
          +-UNICODE+
          '-FSID---'

>-----+-----+-----+-----><
'-TOTime---time-'
```

Parameters

pool_name (Required)

Specifies the name of the directory-container storage pool that is reported in the data deduplication statistics. You can specify up to 30 characters for the storage pool name. If you specify more than 30 characters the command fails.

Restriction: You can only specify directory-container storage pools or cloud storage pools.

node_name

Specifies the name of the client node that is reported in the data deduplication statistics. This parameter is optional. If you do not specify a value for this parameter, all nodes are displayed. You can specify up to 64 characters for the node name. If you specify more than 64 characters the command fails.

file_space_name or FSID

Specifies the name or file space ID (FSID) of one or more file spaces that is reported in the data deduplication statistics. This parameter is optional. You can use wildcard characters to specify this name. An asterisk is the default. Specify one of the following values:

*

Specify an asterisk (*) to show all file spaces or IDs.

file_space_name

Specifies the name of the file space. Specify more than one file space by separating the names with commas and no intervening spaces. FSID Specifies the file space identifier. This parameter is valid for clients with file spaces that are in Unicode format. Specify more than one file space by separating the names with commas and no intervening spaces.

For clients with file spaces that are in Unicode format, you can enter either a file space name or a FSID. If you enter a file space name, the server might have to convert the file space name that you enter. For example, the server might have to convert the name that you enter from the server's code page to Unicode.

Restrictions: The following restrictions apply to file space names and file space identifiers (FSID):

- You must specify a node name if you specify a file space name.
- Do not specify both file space names and FSIDs on the same command.

CODEType

Specifies what type of file spaces to include in the report. The default value is BOTH, which specifies that file spaces are included regardless of code page type. Use this parameter only when you enter an asterisk to display information about all file spaces. This parameter is optional. Specify one of the following values:

UNICODE

Include file spaces that are in Unicode format.

NONUNICODE

Include file spaces that are not in Unicode format.

BOTH

Include file spaces regardless of code page type. This is the default.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. Use this parameter when IBM Spectrum Protect™ clients have file spaces that are in Unicode format, and are on Windows, NetWare, or Macintosh OS X operating systems. This parameter is optional.

This parameter is required if you specify a node name and a file space name or FSID.

Restriction: When you specify this parameter, the file space name cannot contain an asterisk.

Specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names. This is the default.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

TODate

Specifies the latest date for statistics to be deleted. IBM Spectrum Protect deletes only those statistics with a date on or before the date you specify. This parameter is optional.

Specify one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	10/15/2015 If you specify a date, all candidate records that are written on that day (ending at 11:59:59 pm) will be evaluated.
TODAY	The current date.	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To display information that is created until yesterday, you can specify TODATE=TODAY-1 or TODATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include records that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include records that were active on the 10th day of the current month.

TOTime

Specifies that you want to delete data deduplication statistics that are created on or before this time on the specified date. This parameter is optional. The default is the end of the day (23:59:59). Specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified date.	12:30:22
NOW	The current time on the specified date.	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified date.	NOW+03:00 or +03:00. If you issue the DELETE DEDUPSTATS command at 9:00 with TOTIME=NOW+03:00 or TOTIME+=03:00, IBM Spectrum Protect deletes records with a time of 12:00 or earlier on the specified date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified date.	NOW-03:30 or -03:30. If you issue the DELETE DEDUPSTATS command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, IBM Spectrum Protect deletes records with a time of 5:30 or earlier on the specified date.




Example: Delete data deduplication statistics for a file space

Delete data deduplication statistics of a file space that is called /srvr that belongs to a directory-container storage pool, POOL1, that is stored on client node NODE1.

```
delete dedupstats pool1 node1 /srvr
```

Related commands

Table 1. Commands related to DELETE DEDUPSTATS

Command	Description
GENERATE DEDUPSTATS	Generates data deduplication statistics.
   QUERY DEDUPSTATS	Displays data deduplication statistics.

DELETE DEVCLASS (Delete a device class)

Use this command to delete a device class.

To use this command, you must first delete all storage pools that are assigned to the device class and, if necessary, cancel any database export or import processes that are using the device class.

You cannot delete the device class DISK, which is predefined at installation, but you can delete any device classes defined by the IBM Spectrum Protect™ administrator.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DELeTe DEVclass--device_class_name-----<<
```

Parameters

device_class_name (Required)
Specifies the name of the device class to be deleted.









Example: Delete a device class

Delete the device class named MYTAPE. There are no storage pools assigned to the device class.

```
delete devclass mytape
```

Related commands

Table 1. Commands related to DELETE DEVCLASS

Command	Description
DEFINE DEVCLASS	Defines a device class.
  DEFINE DEVCLASS (z/OS® media server)	  Defines a device class to use storage managed by a z/OS media server.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE DEVCLASS	Changes the attributes of a device class.
  UPDATE DEVCLASS (z/OS media server)	  Changes the attributes of a device class for storage managed by a z/OS media server.

DELETE DOMAIN (Delete a policy domain)

Use this command to delete a policy domain. All associated policy sets, including the ACTIVE policy set, management classes, and copy groups are deleted along with the policy domain.

You cannot delete a policy domain to which client nodes are registered. To determine if any client nodes are registered to a policy domain, issue the QUERY DOMAIN or the QUERY NODE command. Move any client nodes to another policy domain, or delete the nodes.

You can delete the predefined STANDARD policy domain. However, if you later reinstall the IBM Spectrum Protect™ server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe D0main--domain_name-----<<
```

Parameters

domain_name (Required)
Specifies the policy domain to delete.

Examples: Delete a policy domain

Delete the EMPLOYEE_RECORDS policy domain.

```
delete domain employee_records
```

Related commands

Table 1. Commands related to DELETE DOMAIN

Command	Description
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
QUERY DOMAIN	Displays information about policy domains.
UPDATE DOMAIN	Changes the attributes of a policy domain.

DELETE DRIVE (Delete a drive from a library)

Use this command to delete a drive from a library. A drive that is in use cannot be deleted.

All paths related to a drive must be deleted before the drive itself can be deleted.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DELeTe DRive--library_name--drive_name-----<<
```

Parameters

library_name (Required)
Specifies the name of the library where the drive is located.

drive_name (Required)
Specifies the name of the drive to be deleted.

Example: Delete a drive from a library

Delete DRIVE3 from the library named AUTO.

```
delete drive auto drive3
```

Related commands

Table 1. Commands related to DELETE DRIVE

Command	Description
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE DRIVE	Changes the attributes of a drive.

DELETE EVENT (Delete event records)

Use this command to delete event records from the database. An event record is created whenever processing of a scheduled command is started or missed.

This command only deletes the event records that exist at the time the command is processed. An event record will not be found:

- If the event record has never been created (the event is scheduled for the future)
- If the event has passed and the event record has already been deleted.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax

```

>>-DElete EVent--date--+-00:00-.----->
                        '-time--'

.-TYPE-----Client-----
>-----<<
'-TYPE-----+Client-----+
      +-Administrative+
      '-All-----'

```

Parameters

date (Required)

Specifies the date used to determine which event records to delete. The maximum number of days you can specify is 9999.

Use this parameter in conjunction with the TIME parameter to specify a date and time for deleting event records. Any record whose scheduled start occurs before the specified date and time is deleted. However, records are not deleted for events whose startup window has not yet passed.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified	TODAY-3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

time

Specifies the time used to determine which event records to delete. Use this parameter in conjunction with the DATE parameter to specify a date and time for deleting event records. Any record whose scheduled start occurs before the specified date and time is deleted. However, records are not deleted for events whose startup window has not yet passed. The default is 00:00.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+03:00 or +03:00 Attention: If you issue this command at 9:00 using NOW+03:00 or +03:00, IBM Spectrum Protect™ deletes records with a time of 12:00 or later on the date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-03:00 or -03:00

TYPE

Specifies the type of events to be deleted. This parameter is optional. The default is CLIENT. Possible values are:

- Client
Specifies to delete event records for client schedules.
- Administrative
Specifies to delete event records for administrative command schedules.
- ALL
Specifies to delete event records for both client and administrative command schedules.

Example: Delete event records

Delete records for events with scheduled start times prior to 08:00 on May 26, 1998 (05/26/1998), and whose startup window has passed. Records for these events are deleted regardless of whether the retention period for event records, as specified with the SET EVENTRETENTION command, has passed.

```
delete event 05/26/1998 08:00
```

Related commands

Table 1. Commands related to DELETE EVENT

Command	Description
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
SET EVENTRETENTION	Specifies the number of days to retain records for scheduled operations.

DELETE EVENTSERVER (Delete the definition of the event server)

Use this command to delete the definition of the event server. You must issue this command before you issue the DELETE SERVER command. If you specify the server defined as the event server on the DELETE SERVER command, you will receive an error message.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe EVENTSErVer-----<<
```

Example: Delete an event server definition

Delete the definition for the event server ASTRO.

```
delete eventserver
```

Related commands

Table 1. Commands related to DELETE EVENTSERVER

Command	Description
DEFINE EVENTSERVER	Defines a server as an event server.
QUERY EVENTSERVER	Displays the name of the event server.

DELETE FILESPACE (Delete client node data from the server)

Use this command to delete file spaces from the server. Files that belong to the file space are deleted from primary, active-data, and copy storage pools, and any file space collocation groups.

IBM Spectrum Protect™ deletes one or more file spaces as a series of batch database transactions, thus preventing a rollback or commit for an entire file space as a single action. If the process is canceled or if a system failure occurs, a partial deletion can occur. A subsequent DELETE FILESPACE command for the same node or owner can delete the remaining data.

If this command is applied to a WORM (write once, read many) volume, the volume is returned to scratch if it has space on which data can be written. (Data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can be written only in space that does not contain current, deleted, or expired data.) If a WORM volume does not have any space available on which data can be written, it remains private. To remove the volume from the library, you must use the CHECKOUT LIBVOLUME command.

Tips:

- If archive retention protection is enabled, the server deletes archive files with expired retention periods. For more information, see the SET ARCHIVERETENTIONPROTECTION command.
- The server does not delete archive files that are on deletion hold until the hold is released.
- Reclamation does not start while the DELETE FILESPACE command is running.
- If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
- If you delete a file space in a deduplicated storage pool, the file space name DELETED is displayed in the output of the QUERY OCCUPANCY command until all deduplication dependencies are removed.
- When replication is configured for a file space, the DELETE FILESPACE command deletes only the file space on the server where you issued the command. If you issue the REPLICATE NODE command, the file space is not deleted on the other replication server.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-DELEte Filespace--node_name--file_space_name----->
      .-Type-----ANY----- .-Data-----ANY----- .
>--+-----+-----+-----+-----+-----+----->
      '-Type-----+ANY-----+' '-Data-----+ANY-----+'
              +-Backup-----+           +-Files-----+
              +-ARchive-----+           |           (1) |
              +-SPacemanaged-+           '-IMages-----'
              '-SERver-----'

      .-Wait-----No----- .
>--+-----+-----+-----+-----+-----+----->
      '-Wait-----+No--+-' '-OWNer-----owner_name-'
              '-Yes-'

      .-NAMEType-----SERVER----- .
>--+-----+-----+-----+-----+-----+----->
      '-NAMEType-----+SERVER--+-'
              +-UNIcode-+
              '-FSID-----'

      .-CODEType-----BOTH----- .
>--+-----+-----+-----+-----+-----+----->>
      '-CODEType-----+UNIcode-----+'
              +-NONUNIcode-+
              '-BOTH-----'
```

Notes:

1. This parameter can be used only when TYPE=ANY or TYPE=BACKUP is specified.

Parameters

node_name (Required)

Specifies the name of the client node to which the file space belongs.

file_space_name (Required)

Specifies the name of the file space to be deleted. This name is case-sensitive and must be entered exactly as it is known to the server. To determine how to enter the name, use the QUERY FILESPACE command. You can use wildcard characters to specify this name.

For a server that has clients with support for Unicode, you might have the server convert the file space name that you enter. For example, you might want to have the server convert the name that you entered from the server's code page, to Unicode. See the NAMETYPE parameter for details. If you do not specify a file space name, or specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

Type

Specifies the type of data to be deleted. This parameter is optional. The default value is ANY. You can use the following values:

ANY

Delete only backed-up versions of files and archived copies of files.

If you specify `delete filespace node_name * type=any`, all backed-up data and archived data in all file spaces for that node are deleted. File spaces are deleted only if they do not contain files that are moved from an IBM Spectrum Protect for Space Management client.

Backup

Delete backup data for the file space.

ARchive

Delete all archived data on the server for the file space.

SPacemanaged

Delete files that are migrated from a user's local file system by an IBM Spectrum Protect for Space Management client. The OWNER parameter is ignored when you specify TYPE=SPACEMANAGED.

SERver

Delete all archived files in all file spaces for a node that is registered as TYPE=SERVER.

DAta

Specifies objects to delete. This parameter is optional. The default value is ANY. You can specify one of the following values:

ANY

Delete files, directories, and images.

FIles

Delete files and directories.

IMages

Delete image objects. You can use this parameter only if you specified TYPE=ANY or TYPE=BACKUP.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. You can specify one of the following values:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before you continue with other tasks. The server displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

OWNer

Restricts the data that is deleted to files that belong to the owner. This parameter is optional; it is ignored when TYPE=SPACEMANAGED. This parameter applies to only multiuser client systems such as AIX®, Linux, and Solaris OS.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. A backup-archive client with support for Unicode is available only for the following operating systems: Windows, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include file spaces that are in Unicode.

NONUNICODE

Include file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

Delete a file space

Delete the C_Drive file space that belongs to the client node HTANG.

```
delete filesystem htang C_Drive
```

Delete all space-managed files for a client node

Delete all files that are migrated from client node APOLLO (that is, all space-managed files).

```
delete filesystem apollo * type=spacemanaged
```

Related commands

Table 1. Commands related to DELETE FILESPACE

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY OCCUPANCY	Displays file space information by storage pool.
QUERY PROCESS	Displays information about background processes.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
RENAME FILESPACE	Renames a client filesystem on the server.

DELETE GRPMEMBER (Delete a server from a server group)

Use this command to delete a server or server group from a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte GRPMEMber--group_name-----member_name+-----><
```

Parameters

group_name (Required)

Specifies the group.

member_name (Required)

Specifies the server or group to delete from the group. To specify multiple names, separate the names with commas and no intervening spaces.

Example: Delete a server from a server group

Delete member PHOENIX from group WEST_COMPLEX.

```
delete grpmember west_complex phoenix
```

Related commands

Table 1. Commands related to DELETE GRPMEMBER

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVER	Deletes the definition of a server.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DELETE LIBRARY (Delete a library)

Use this command to delete a library. Before you delete a library, you must delete other associated objects, such as the path.

Use this command to delete a library. Before you delete a library, delete the path and all associated drives.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DELEte LIBRary--library_name-----><
```

Parameters

library_name (Required)
Specifies the name of the library to be deleted.

Example: Delete a manual library

Delete the manual library named LIBR1.

```
delete library libr1
```

Related commands

Table 1. Commands related to DELETE LIBRARY

Command	Description
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE PATH	Changes the attributes associated with a path.

DELETE MACHINE (Delete machine information)

Use this command to delete machine description information. To replace existing information, issue this command and then issue an INSERT MACHINE command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte MACHine--machine_name----->
      .-Type----All-----
>--+-----+----->>
      '-Type----+-All-----+'
          +-RECOVERYInstructions+
          '-Characteristics-----'
```

Parameters

machine_name (Required)
Specifies the name of the machine whose information is to be deleted.

Type
Specifies the type of machine information. This parameter is optional. The default is ALL. Possible values are:

All

- Specifies all information.
- RECOVERYInstructions
 - Specifies the recovery instructions.
- CCharacteristics
 - Specifies the machine characteristics.

Example: Delete a specific machine's information

Delete the machine characteristics associated with the DISTRICT5 machine.

```
delete machine district5 type=characteristics
```

Related commands

Table 1. Commands related to DELETE MACHINE

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the IBM Spectrum ProtectIBM Spectrum Protect™ database.
QUERY MACHINE	Displays information about machines.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE MACHINE	Changes the information for a machine.

DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)

Use this command to delete the association between a machine and one or more nodes. This command does not delete the node from IBM Spectrum Protect™.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-------.
      v           |
>>>-DELEte MACHNODEAssociation--machine_name----node_name-+----->>>

```

Parameters

machine_name (Required)

Specifies the name of a machine that is associated with one or more nodes.

node_name (Required)

Specifies the name of a node associated with a machine. If you specify a list of node names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name. If a node is not associated with the machine, that node is ignored.

Example: Delete an association between a node and a machine

Delete the association between the DISTRICT5 machine and the ACCOUNTSPAYABLE node.

```
delete machnodeassociation district5 accountspayable
```

Related commands

Table 1. Commands related to DELETE MACHNODEASSOCIATION

Command	Description
DEFINE MACHNODEASSOCIATION	Associates an IBM Spectrum Protect node with a machine.
QUERY MACHINE	Displays information about machines.

DELETE MGMTCLASS (Delete a management class)

Use this command to delete a management class. You cannot delete a management class in the ACTIVE policy set. All copy groups in the management class are deleted along with the management class.

You can delete the management class assigned as the default for a policy set, but a policy set cannot be activated unless it has a default management class.

You can delete the predefined STANDARD management class in the STANDARD policy domain. However, if you later reinstall the IBM Spectrum Protect™ server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax

```
>>-DELeTe MGmtclass--domain_name--policy_set_name--class_name--<<
```

Parameters

- domain_name (Required)
Specifies the policy domain to which the management class belongs.
- policy_set_name (Required)
Specifies the policy set to which the management class belongs.
- class_name (Required)
Specifies the management class to delete.

Example: Delete a management class

Delete the ACTIVEFILES management class from the VACATION policy set of the EMPLOYEE_RECORDS policy domain.

```
delete mgmtclass employee_records  
vacation activefiles
```

Related commands

Table 1. Commands related to DELETE MGMTCLASS

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
QUERY MGMTCLASS	Displays information about management classes.
UPDATE MGMTCLASS	Changes the attributes of a management class.

DELETE NODEGROUP (Delete a node group)

Use this command to delete a node group. You cannot delete a node group if it has any members in it.

Attention: You can remove all the members in the node group by issuing the DELETE NODEGROUPMEMBER command with a wildcard in the node_name parameter.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
>>-DELEte NODEGrouP--group_name-----<<
```

Parameters

group_name
Specifies the name of the node group that you want to delete.

Example: Delete a node group

Delete a node group named group1.

```
delete nodegroup group1
```

Related commands

Table 1. Commands related to DELETE NODEGROUP

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DELETE NODEGROUPMEMBER (Delete node group member)

Use this command to delete a client node from a node group.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
>>-DELEte NODEGROUPEmber--group_name----node_name+-----<<
```


Parameters

group_name

Specifies the name of the node group from which you want to delete a client node.

node_name

Specifies the name of the client node that you want to delete from the node group. You can specify one or more names. When specifying multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple nodes.

Example: Delete node group members

Delete two nodes, `node1` and `node2`, from a node group, `group1`.

```
delete nodegroupmember group1 node1,node2
```

Related commands

Table 1. Commands related to DELETE NODEGROUPMEMBER

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DELETE PATH (Delete a path)

Use this command to delete a path definition

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DELEte PATH--source_name--destination_name----->
                                     (1)
>--SRCType-----+DATAMover-----+----->
                   '-SERVer-----'

                                     (2)
>--DESTType-----+DRive-----LIBRARY----library_name+-----<
                   '-LIBRARY-----'
```

Notes:

1. This parameter is only available on AIX, HP-UX, Linux, Solaris, Windows operating systems.
2. This parameter is only available on AIX, HP-UX, Linux, Solaris, Windows operating systems.

Parameters

source_name (Required)

Specifies the name of the source of the path to be deleted. This parameter is required.

The name specified must be that of a server or data mover that is already defined to the server.

destination_name (Required)

Specifies the name of the destination of the path to be deleted. This parameter is required.

SRCType (Required)

Specifies the source type of the path to be deleted. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVER

Specifies that a storage agent is the source.

DESTType (Required)

Specifies the type of the destination. Possible values are:

DRive LIBRARY=library_name

Specifies that a drive is the destination. The DRIVE and LIBRARY parameters are both required when the destination type is drive.

LIBRARY

Specifies that a library is the destination.

Attention: If the path from a data mover to a library is deleted, or the path from the server to a library is deleted, the server will not be able to access the library. If the server is halted and restarted while in this state, the library will not be initialized.

Example: Delete a NAS data mover path

Delete a path from a NAS data mover NAS1 to the library NASLIB.

```
delete path nas1 naslib srctype=datamover desttype=library
```

Related commands

Table 1. Commands related to DELETE PATH

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE PATH	Defines a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE PATH	Changes the attributes associated with a path.

DELETE POLICYSET (Delete a policy set)

Use this command to delete a policy set. When you delete a policy set, all management classes and copy groups that belong to the policy set are also deleted.

The ACTIVE policy set in a policy domain cannot be deleted. You can replace the contents of the ACTIVE policy set by activating a different policy set. Otherwise, the only way to remove the ACTIVE policy set is to delete the policy domain that contains the policy set.

You can delete the predefined STANDARD policy set. However, if you later reinstall the IBM Spectrum Protect™ server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-DELEte Policyset--domain_name--policy_set_name-----><
```

Parameters

domain_name (Required)
Specifies the policy domain to which the policy set belongs.

policy_set_name (Required)
Specifies the policy set to delete.

Example: Delete a policy set

Delete the VACATION policy set from the EMPLOYEE_RECORDS policy domain by issuing the following command:

```
delete policyset employee_records vacation
```

Related commands

Table 1. Commands related to DELETE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

DELETE PROFASSOCIATION (Delete a profile association)

Use this command on a configuration manager to delete the association of one or more objects from a profile. If associations are deleted, the objects are no longer distributed to subscribing managed servers. When managed servers request updated configuration information, the configuration manager notifies them of the object deletions.

A managed server deletes the objects that were deleted from the profile, unless the objects are associated with another profile to which that server subscribes.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte PROFASSOCIation--profile_name----->  
>--+-----+----->  
  '-ADMinS---+*-----+-'  
      | .,----- . |  
      | V           | |  
      '---admin_name-+-'  
>--+-----+----->  
  '-DObains---+*-----+-'
```

```

| .-,----- . |
| V           | |
|'---domain_name+--'

>--+-----+----->
|'-ADSHeds---+*-----+-'
| .-,----- . |
| V           | |
|'---schedule_name+--'

>--+-----+----->
|'-SCRipts---+*-----+-'
| .-,----- . |
| V           | |
|'---script_name+--'

>--+-----+----->
|'-CLOptsets---+*-----+-'
| .-,----- . |
| V           | |
|'---option_set_name+--'

>--+-----+----->
|'-SERVers---+*-----+-'
| .-,----- . |
| V           | |
|'---server_name+--'

>--+-----+-----><
|'-SERVERGroups---+*-----+-'
| .-,----- . |
| V           | |
|'---group_name+--'

```

Parameters

profile_name (Required)

Specifies the profile from which to delete associations.

ADMins

Specifies the administrators whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all administrators from the profile. If you specify a list of administrators and a match-all definition exists for the profile, the command fails. Administrator definitions are not changed on the configuration manager. However, they are automatically deleted from all subscribing managed servers at the next configuration refresh, with the following exceptions:

- An administrator is not deleted if that administrator has an open session on the server.
- An administrator is not deleted if, as a result, the managed server would have no administrators with system privilege class.

DOmains

Specifies the domains whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all domains from the profile. If you specify a list of domains and a match-all domain definition exists for the profile, the command fails.

The domain information is automatically deleted from all subscribing managed servers. However, a policy domain that has client nodes assigned will not be deleted. To delete the domain at the managed server, assign those client nodes to another policy domain.

ADSHeds

Specifies a list of administrative schedules whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. If you specify a list of administrative schedules and a match-all administrative schedule definition exists for the profile, the command fails. Use the match-all character (*) to delete all administrative schedules from the profile.

The administrative schedules are automatically deleted from all subscribing managed servers. However, an administrative schedule is not deleted if the schedule is active on the managed server. To delete an active schedule, make the schedule inactive.

SCRipts

Specifies the server command scripts whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all scripts from the profile. If you specify a list of scripts and a match-all script definition exists for the profile, the command fails. The server command scripts are automatically deleted from all subscribing managed servers.

CLOptsets

Specifies the client option sets whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all client option sets from the profile. If you specify a list of client option sets and a match-all client option set definition exists for the profile, the command fails. The client option sets are automatically deleted from all subscribing managed servers.

SERVers

Specifies the servers whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. You can use the match-all character (*) to delete all servers from the profile. If you specify a list of servers and a match-all server definition exists for the profile, the command fails. The server definitions are automatically deleted from all subscribing managed servers with the following exceptions:

- A server definition is not deleted if the managed server has an open connection to another server.
- A server definition is not deleted if the managed server has a device class of the device type SERVER that refers to the other server.
- A server definition is not deleted if the server is the event server for the managed server.

SERVERGroups

Specifies the server groups whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. You can use the match-all character (*) to delete all server groups from the profile. If you specify a list of server groups and a match-all group definition exists for the profile, the command fails. The server group definitions are automatically deleted from all subscribing managed servers.

Example: Delete the domain associations for a specific profile

Delete all domain associations from a profile named MIKE.

```
delete profassociation mike domains=*
```

Related commands

Table 1. Commands related to DELETE PROFASSOCIATION

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DELETE PROFILE (Delete a profile)

Use this command on a configuration manager to delete a profile and stop its distribution to managed servers.

You cannot delete a locked profile. You must first unlock the profile with the UNLOCK PROFILE command.

Deleting a profile from a configuration manager does not delete objects associated with that profile from the managed servers. You can use the DELETE SUBSCRIPTION command with the DISCARDOBJECTS=YES parameter on each subscribing managed

server to delete subscriptions to the profile and associated objects. This also prevents the managed servers from requesting further updates to the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte PROFILE--profile_name--+-Force-----No-----+-----><
                                     '-Force-----+No--+-'
                                     '-Yes-'
```

Parameters

profile_name (Required)

Specifies the profile to delete.

Force

Specifies whether the profile is deleted if one or more managed servers have subscriptions to that profile. The default is NO. Possible values are:

No

Specifies that the profile is not deleted if one or more managed servers have subscriptions to that profile. You can delete the subscriptions on each managed server using the DELETE SUBSCRIPTION command.

Yes

Specifies that the profile is deleted even if one or more managed servers have subscriptions to that profile. Each subscribing server continues to request updates for the deleted profile until the subscription is deleted.

Examples: Delete a profile

Delete a profile named BETA, even if one or more managed servers subscribe to it.

```
delete profile beta force=yes
```

Related commands

Table 1. Commands related to DELETE PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)

Use this command to remove the association of one or more machines with a recovery media. This command does not delete the machine from IBM Spectrum Protect™.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
                .-----  
                |  
>>-DELEte RECMEDMACHAssociation--media_name----machine_name-+---><
```

Parameters

media_name (Required)

Specifies the name of the recovery media that is associated with one or more machines.

machine_name (Required)

Specifies the name of the machine associated with the recovery media. To specify a list of machine names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name. If a machine is not associated with the recovery media, the machine is ignored.

Example: Delete a machine's association with recovery media

Delete the association between the DIST5RM recovery media and the DISTRICT1 and DISTRICT5 machines.

```
delete recmedmachassociation  
dist5rm district1,district5
```

Related commands

Table 1. Commands related to DELETE RECMEDMACHASSOCIATION

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
QUERY MACHINE	Displays information about machines.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.

DELETE RECOVERYMEDIA (Delete recovery media)

Use this command to delete a recovery media definition from IBM Spectrum Protect™.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte RECOVERYMedia--media_name-----><
```

Parameters

media_name (Required)
Specifies the name of the recovery media.

Example: Delete a recovery media definition

Delete the DIST5RM recovery media.

```
delete recoverymedia dist5rm
```

Related commands

Table 1. Commands related to DELETE RECOVERYMEDIA

Command	Description
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

DELETE SCHEDULE (Delete a client or an administrative command schedule)

Use this command to delete schedules from the database.

The DELETE SCHEDULE command takes two forms: one if the schedule applies to client operations, one if the schedule applies to administrative commands. The syntax and parameters for each form are defined separately.

Table 1. Commands related to DELETE SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
QUERY SCHEDULE	Displays information about schedules.
UPDATE SCHEDULE	Changes the attributes of a schedule.

- DELETE SCHEDULE (Delete a client schedule)
Use the DELETE SCHEDULE command to delete one or more client schedules from the database. Any client associations to a schedule are removed when the schedule is deleted.
- DELETE SCHEDULE (Delete an administrative schedule)
Use this command to delete one or more administrative command schedules from the database.

DELETE SCHEDULE (Delete a client schedule)

Use the DELETE SCHEDULE command to delete one or more client schedules from the database. Any client associations to a schedule are removed when the schedule is deleted.

Privilege class

To delete a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the specified policy domain.

Syntax

```
>>-DELeTe SChedule--domain_name--schedule_name----->  
. -Type-----Client-
```


>--+-----+-----><

Parameters

- domain_name (Required)
Specifies the name of the policy domain to which the schedule belongs.
- schedule_name (Required)
Specifies the name of the schedule to delete. You can use a wildcard character to specify this name.
- Type=Client
Specifies to delete a client schedule. This parameter is optional. The default is CLIENT.

Example: Delete a specific schedule from a specific policy domain

Delete the WEEKLY_BACKUP schedule, which belongs to the EMPLOYEE_RECORDS policy domain.

```
delete schedule employee_records weekly_backup
```

DELETE SCHEDULE (Delete an administrative schedule)

Use this command to delete one or more administrative command schedules from the database.

Privilege class

To delete an administrative command schedule, you must have system authority.

Syntax

```
>>-DELEte SCHedule--schedule_name--Type--===Administrative-----><
```

Parameters

- schedule_name (Required)
Specifies the name of the schedule to delete. You can use a wildcard character to specify this name.
- Type=Administrative (Required)
Specifies to delete an administrative command schedule.

Example: Delete an administrative command schedule

Delete the administrative command scheduled named DATA_ENG.

```
delete schedule data_eng type=administrative
```

DELETE SCRATCHPADENTRY (Delete a scratch pad entry)

Use this command to delete one or more lines of data from a scratch pad.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte SCRATCHPadentry--major_category--minor_category----->  
      .-Line---*-----.  
>--subject--+-----+-----><  
      '-Line---number-'
```

Parameters

major_category (Required)

Specifies the major category from which one or more lines of data are to be deleted. This parameter is case sensitive.

minor_category (Required)

Specifies the minor category from which one or more lines of data are to be deleted. This parameter is case sensitive.

subject (Required)

Specifies the subject from which one or more lines of data are to be deleted. This parameter is case sensitive.

Line

Specifies a line of data that is to be deleted. For number, enter the number of the line that is to be deleted. All data on the line is deleted. The numbering of other lines in the subject section is not affected. You can delete all lines of data from a subject section by omitting the Line parameter in this command.

Example: Delete all lines of data from a subject in a scratch pad

Delete all lines of data about the location of an administrator, Jane, from a database that stores information about administrators:

```
delete scratchpadentry admin_info location jane
```

Related commands

Table 1. Commands related to DELETE SCRATCHPADENTRY

Command	Description
DEFINE SCRATCHPADENTRY	Creates a line of data in the scratch pad.
QUERY SCRATCHPADENTRY	Displays information that is contained in the scratch pad.
SET SCRATCHPADRETENTION	Specifies the amount of time for which scratch pad entries are retained.
UPDATE SCRATCHPADENTRY	Updates data on a line in the scratch pad.

DELETE SCRIPT (Delete command lines from a script or delete the entire script)

Use this command to delete a single line from an IBM Spectrum Protect™ script or to delete the entire IBM Spectrum Protect script.

Privilege class

To issue this command, the administrator must have previously defined the script or must have system privilege.

Syntax

```
>>-DELEte SCRipt--script_name--+-----+-----><  
                                '-Line----number-'
```

Parameters

script_name (Required)

Specifies the name of the script to delete. The script is deleted unless you specify a line number.

Line

Specifies the line number to delete from the script. If you do not specify a line number, the entire script is deleted.

Example: Delete a specific line from a script

Using the following script named QSAMPLE and issue a command to delete line 005 from it.

```

001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS

delete script qsample line=5

```

Related commands

Table 1. Commands related to DELETE SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

DELETE SERVER (Delete a server definition)

Use this command to delete a server definition.

This command fails if the server:

- Is defined as the event server.
- Is named in a device class definition whose device type is SERVER.
- Has an open connection to or from another server.
- Is a target server for virtual volumes.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte--SERver--server_name-----><
```

Parameters

server_name (Required)
Specifies a server name.

Example: Delete a server's definition

Delete the definition for a server named SERVER2.

```
delete server server2
```

Related commands

Table 1. Commands related to DELETE SERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
QUERY EVENTSERVER	Displays the name of the event server.
QUERY SERVER	Displays information about servers.

Command	Description
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
UPDATE SERVER	Updates information about a server.

DELETE SERVERGROUP (Delete a server group)

Use this command to delete a server group. If the group you delete is a member of other server groups, IBM Spectrum Protect™ also removes the group from the other groups.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe SERVERGroup--group_name-----><
```

Parameters

group_name (Required)
Specifies the server group to delete.

Example: Delete a server group

Delete a server group named WEST_COMPLEX.

```
delete servergroup west_complex
```

Related commands

Table 1. Commands related to DELETE SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DELETE SPACETRIGGER (Delete the storage pool space triggers)

Use this command to delete the definition of the storage pool space trigger.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-DELeTe SPACETriGger--STG----->
>--+-----+-----><
  '-STGPOOL---storage_pool_name-'
```

Parameters

STG

Specifies a storage pool space trigger.

STGPOOL

Specifies the storage pool trigger to be deleted. If STG is specified without specifying STGPOOL, the default storage pool space trigger is the deletion target.

Example: Delete a space trigger definition

Delete the space trigger definition for the WINPOOL1 storage pool.

```
delete spacetrigger stg stgpool=winpool1
```

Related commands

Table 1. Commands related to DELETE SPACETRIGGER

Command	Description
DEFINE SPACETRIGGER	Defines a space trigger to expand the space for a storage pool.
QUERY SPACETRIGGER	Displays information about a storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)

Use this command to delete an existing status monitoring threshold.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe STATusthreshoLd--threshold_name-----><
```

Parameters

threshold_name (Required)

Specifies the threshold name that you want to delete.

Delete an existing status threshold

Delete an existing status threshold by issuing the following command:

Related commands

Table 1. Commands related to DELETE STATUSTHRESHOLD

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

DELETE STGPOOL (Delete a storage pool)

Use this command to delete a storage pool. To delete a storage pool, you must first delete all volumes that are assigned to the storage pool.

You cannot delete a storage pool that is identified as the next storage pool for another storage pool. For more information about storage pool hierarchy, see the NEXTSTGPOOL parameter in the DEFINE STGPOOL command.

Restrictions:

- For container storage pools, delete all storage pool directories before you delete the storage pool.
- Do not delete a storage pool that is specified as a destination for a management class or copy group in the ACTIVE policy set. Client operations might fail as a result.
- When you delete a copy storage pool that was previously included in a primary storage-pool definition (specifically in the COPYSTGPOOLS list), you must remove the copy storage pool from the list before deletion. Otherwise, the DELETE STGPOOL command fails until all references to that copy pool are removed. For each primary storage pool with a reference to the copy storage pool to be deleted, remove the reference by entering the UPDATE STGPOOL command with the COPYSTGPOOLS parameter with all previous copy storage pools except the copy storage pool to be deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe STGpool--pool_name-----<<
```

Parameters

pool_name (Required)
Specifies the storage pool to delete.

Example: Delete a storage pool

Delete the storage pool named POOLA.

delete stgpool poola

Related commands

Table 1. Commands related to DELETE STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
DELETE STGPOOLDIRECTORY	Deletes a storage pool directory from a directory-container or cloud-container storage pool.
QUERY STGPOOL	Displays information about storage pools.
QUERY STGPOOLDIRECTORY	Displays information about storage pool directories.
AIX Windows SET DRMCOPYSTGPOOL	AIX Windows Specifies that copy storage pools are managed by DRM.
UPDATE STGPOOL	Changes the attributes of a storage pool.
UPDATE STGPOOLDIRECTORY	Changes the attributes of a storage pool directory.

DELETE STGPOOLDIRECTORY (Deleting a storage pool directory)

Use this command to delete a definition for a storage pool directory.

You might want to delete a storage pool directory for the following reasons:

- To decommission old storage.
- To discontinue using the local disk before moving data to the cloud.
- To no longer maintain the data in the storage pool directory because there is no requirement to do so.

Restrictions:

- You can issue this command only when no containers are assigned to the storage pool directory. Issue the QUERY CONTAINER command to determine whether any containers are assigned to the storage pool directory.
- To remove containers from a storage pool directory, you must issue the UPDATE STGPOOLDIRECTORY command and specify the ACCESS=DESTROYED parameter. Then, issue the AUDIT CONTAINER command and specify the ACTION=REMOVEDAMAGED parameter. Verify that the containers are removed. The ACTION=REMOVEDAMAGED parameter removes the inventory information of the objects that were backed up or archived. You should only remove the inventory information if you do not need the backups.

If you experience a hardware failure or a loss of your directory, see the relevant AUDIT and REPAIR commands. You should make any repairs to the IBM Spectrum Protect™ environment before you delete the storage pool directory.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe STGPOOLDIRectory--pool_name--directory-----><
```

Parameters

pool_name (Required)

Specifies the storage pool that contains the directory to delete. This parameter is required.

directory (Required)

Specifies the file system directory of the storage pool to delete. This parameter is required.

Example: Update a storage pool directory to prepare for deletion

Update the storage pool directory that is named DIR1 in storage pool POOLA to mark as destroyed. When a storage pool is marked as destroyed, you can delete it.

AIX Linux

```
update stgpooldirectory poola /storage/dir1 access=destroyed
```

Windows

```
update stgpooldirectory poola e:\storage\dir1 access=destroyed
```

Example: Delete a storage pool directory

Delete the storage pool directory that is named DIR1 in storage pool POOLA.

AIX Linux

```
delete stgpooldirectory poola /storage/dir1
```

Windows

```
delete stgpooldirectory poola e:\storage\dir1
```

Table 1. Commands related to DELETE STGPOOLDIRECTORY

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
QUERY STGPOOLDIRECTORY	Displays information about storage pool directories.
UPDATE STGPOOLDIRECTORY	Changes the attributes of a storage pool directory.
QUERY EXTENTUPDATES	Displays information about updates to data extents in directory-container storage pools.

DELETE STGRULE (Delete storage rules for storage pools)

Use this command to delete storage rules for one or more storage pools.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DElete STGRULE--rule_name-----><
```

Parameters

rule_name(Required)

Specifies the name of the storage rule that must be deleted. The maximum length of the name is 30 characters.

Delete a storage rule

Delete a storage rule that is named stgrule1:


```
delete stgrule stgrule1
```

Related commands

Table 1. Commands related to DELETE STGRULE

Command	Description
DEFINE STGRULE (tiering)	Defines a storage rule for tiering.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (tiering)	Updates a tiering storage rule.

DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)

Use this command on a configuration manager to delete managed server subscriptions from the configuration manager database. Use this command when a managed server no longer exists or cannot notify the configuration manager after deleting a subscription.

Attention: Use this command only in rare situations in which the configuration manager's database contains an entry for a subscription, but the managed server does not have such a subscription. For example, use this command if a managed server no longer exists or cannot notify the configuration manager after deleting a subscription.

Under normal circumstances, use the DELETE SUBSCRIPTION command to delete a subscription from the managed server. The managed server notifies the configuration manager, which then deletes the subscription from its database.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELeTe SUBSCRIBer--server_name-----<<
```

Parameters

server_name (Required)

Specifies the name of the managed server with subscription entries to be deleted.

Example: Delete subscription entries for a specific managed server

Delete all subscription entries for a managed server named DAN.

```
delete subscriber dan
```

Related commands

Table 1. Commands related to DELETE SUBSCRIBER

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.

DELETE SUBSCRIPTION (Delete a profile subscription)

Use this command on a managed server to delete a profile subscription. You can also delete from the managed server all objects associated with the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DELEte SUBSCRIPtion--profile_name----->
. -DISCARDobjects----No-----
>--+-----+-----+----->>
' -DISCARDobjects----+-No--+-'
      '-Yes-'
```

Parameters

profile_name (Required)

Specifies the name of the profile for which the subscription is to be deleted.

DISCARDobjects

Specifies whether objects associated with the profile are to be deleted on the managed server. This parameter is optional. The default is NO.

No

Specifies that the objects are not to be deleted.

Yes

Specifies that the objects are to be deleted, unless they are associated with another profile for which a subscription is defined.

Example: Delete a profile subscription

Delete a subscription to a profile named ALPHA and its associated objects from a managed server.

```
delete subscription alpha discardobjects=yes
```

Related commands

Table 1. Commands related to DELETE SUBSCRIPTION

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.

DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)

Use this command to delete a virtual file space mapping definition. Virtual file spaces containing data cannot be deleted unless you use the DELETE FILESPACE command first.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned

Syntax

```
>>-DElete VIRTUALFSmapping -node_name----->  
>--virtual_filespace_name-----<<
```

Parameters

node_name (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

virtual_filespace_name (Required)

Specifies the name of the virtual file space mapping definition to be deleted. Wildcard characters are allowed.

Example: Delete a virtual file space mapping

Delete the virtual file space mapping definition /mikeshomedir for the NAS node named NAS1.

```
delete virtualfsmapping nas1 /mikeshomedir
```

Related commands

Table 1. Commands related to DELETE VIRTUALFSMAPPING

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY VIRTUALFSMAPPING	Query a virtual file space mapping.
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

DELETE VOLHISTORY (Delete sequential volume history information)

Use this command to delete volume history file records that are no longer needed (for example, records for obsolete database backup volumes).

When you delete records for volumes that are not in storage pools (for example, database backup or export volumes), the volumes return to scratch status even if IBM Spectrum Protect™ acquired them as private volumes. Scratch volumes of device type FILE are deleted. When you delete the records for storage pool volumes, the volumes remain in the IBM Spectrum Protect database. When you delete records for recovery plan file objects from a source server, the objects on the target server are marked for deletion.

Restriction: Do not use the DELETE VOLHISTORY command to delete information about backup set volumes from the volume history file. Instead, use the DELETE BACKUPSET command for this purpose.

For users of DRM, the database backup expiration should be controlled with the SET DRMDBBACKUPEXPIREDAYS command instead of this DELETE VOLHISTORY command. Use the DELETE VOLHISTORY command to remove a record of the volume. This can cause volumes to be lost that were managed by the MOVE DRMEDIA command. Use the SET DRMDBBACKUPEXPIREDAYS command to manage the automatic expiration of DRM database backup volumes.

Tips:

- Volumes for the most recent database backup series are not deleted.
- Existing volume history files are not automatically updated with this command.
- You can use the DEFINE SCHEDULE command to periodically delete volume history records.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DElete VOLHistory--TODate---date----->
. -Totime---23:59:59-
>--+-----+----->
'-Totime---time-----'

>--Type---+All-----><
+-DBBackup--+-----+
|           '-DEVclass---class_name-' |
+-DBSnapshot--+-----+
|           '-DEVclass---class_name-' |
+-DBRpf-----+
+-EXPort-----+
|           .-DELETEDatest---No----- |
+-RPFile--+-----+
|           '-DELETEDatest---+No--+-' |
|           '-Yes-' |
|           .-DELETEDatest---No----- |
+-RPFSnapshot--+-----+
|           '-DELETEDatest---+No--+-' |
|           '-Yes-' |
+-STGNew-----+
+-STGReuse-----+
'-STGDelete-----'
```

Parameters

TODate (Required)

Specifies the date to use to select sequential volume history information to be deleted. You can delete only those records with a date on or before the date that you specify. You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	01/23/1999
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-30 or -30. To delete records that are 30 or more days old, you can specify TODAY-30 or simply -30.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

TOTime

Specifies that you want to delete records that are created on or before this time on the specified date. This parameter is optional. The default is the end of the day (23:59:59). You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified date	12:30:22
NOW	The current time on the specified date	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified date	NOW+03:00 or +03:00. If you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW+03:00 or TOTIME=+03:00, IBM Spectrum Protect deletes records with a time of 12:00 or earlier on the specified date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified date	NOW-03:30 or -03:30. If you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, IBM Spectrum Protect deletes records with a time of 5:30 or earlier on the specified date.

Type (Required)

Specifies the type of records, which also meet the date and time criteria, to delete from the volume history file. Possible values are:

All

Specifies to delete all records.

Restriction: The DELETE VOLHISTORY command does not delete records of remote volumes.

DBBackup

Specifies to delete only records that contain information about volumes that are used for database full and incremental backups, that is, with volume types of BACKUPFULL and BACKUPINCR, and that meet the specified date and time criteria. The records from the latest full and incremental database backup series will not be deleted.

DEVclass=class_name

Specifies the device class name that was used to create the database backups. This optional parameter can be used to delete database backups that are created by using a server-to-server virtual volume device class. The type of the device class must be SERVER. This parameter can be used only to delete volume history entries of type BACKUPFULL, BACKUPINCR, or DBSNAPSHOT.

A full or incremental database backup volume is eligible to be deleted if all of the following conditions are met:

- The device class that was used to create the database backup volume matches the specified device class.
- The volume was created on or before the specified date and time.
- The volume is not part of the latest full plus incremental database backup series.
- The volume is not part of a full plus incremental backup series with an incremental database backup that was created after the specified date and time.

DBSnapshot

Specifies to delete only records that contain information about volumes that are used for snapshot database backups, and that meet the specified date and time criteria. Records that are related to the latest snapshot database backup will not be deleted.

DEVclass=classname

Specifies the device class name that was used to create the database backups. This optional parameter can be used to delete database backups that are created by using a server-to-server virtual volume device class. The type of the device class must be SERVER. This parameter can only be used to delete volume history entries of type BACKUPFULL, BACKUPINCR, or DBSNAPSHOT.

A snapshot database backup volume is eligible to be deleted if all of the following conditions are met:

- The device class that is used to create the database backup volume matches the specified device class
- The volume was created on or before the specified date and time
- The volume is not part of the latest snapshot database backup series

DBRpf

Specifies to delete only records that contain information about full and incremental database backup volumes and recovery plan file volumes.

EXPort

Specifies to delete only records that contain information about export volumes.

RPFfile

Specifies to delete only records that contain information about recovery plan file objects that are stored on a target server and that meet the specified date and time criteria.

DELETEDlatest

Specifies whether the latest recovery plan file is eligible for deletion. This optional parameter can be used to delete the latest recovery plan files that are created by using a server-to-server virtual volume device class.

This parameter can be used only to delete volume history entries of type RPFfile (for instance, those recovery plan files that were created by using the DEVCLASS parameter with the PREPARE command). If this parameter is not specified, the latest RPFfile entries are not deleted.

No

Specifies the latest RPFfile file is not deleted.

Yes

Specifies the latest RPFfile file is deleted if it meets the specified date and time criteria.

RPFSSnapshot

Specifies to delete only records that contain information about recovery plan file objects that were created for snapshot database backups, that are stored on a target server and that meet the specified date and time criteria. The latest RPFSSNAPSHOT file will not be deleted unless it meets the specified date and time criteria, and the DELETE parameter is set to Yes.

DELETEDlatest

Specifies whether the latest recovery plan file is eligible for deletion. This optional parameter can be used to delete the latest recovery plan files that are created by using a server-to-server virtual volume device class.

This parameter can only be used to delete volume history entries of type RPFSSNAPSHOT (for instance, those recovery plan files that were created by using the DEVCLASS parameter with the PREPARE command). If this parameter is not specified, the latest RPFSSNAPSHOT entries are not deleted.

No

Specifies the latest RPFSSNAPSHOT file is not deleted.

Yes

Specifies the latest RPFSSNAPSHOT file is deleted if it meets the specified date and time criteria.

STGNew

Specifies to delete only records that contain information about new sequential access storage volumes.

STGReuse

Specifies to delete only records that contain information about reused sequential storage pool volumes.

STGDelete

Specifies to delete only records that contain information about deleted sequential storage pool volumes.

Example: Delete recovery plan file information

Delete all recovery plan file information that is created on or before 03/28/2016.

```
delete volhistory type=rpfile todate=03/28/2016
```

Related commands

Table 1. Commands related to DELETE VOLHISTORY

Command	Description
BACKUP VOLHISTORY	Records volume history information in external files.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE VOLUME	Deletes a volume from a storage pool.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
MOVE DRMEDIA	Moves DRM media onsite and offsite.

Command	Description
PREPARE	Creates a recovery plan file.
QUERY RPFIL	Displays information about recovery plan files.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMRPFEXPIREDAYS	Set criteria for recovery plan file expiration.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.

DELETE VOLUME (Delete a storage pool volume)

Use this command to delete a storage pool volume and, optionally, the files stored in the volume.

If the volume has data, to delete the volume you must do one of the following:

- Before deleting the volume, use the MOVE DATA command to move all files to another volume.
- Explicitly request to discard all files in the volume when the volume is deleted (by specifying DISCARDATA=YES).

If you are deleting several volumes, delete the volumes one at a time. Deleting more than one volume at a time can adversely affect server performance.

Storage pool volumes cannot be deleted if they are in use. For example, a volume cannot be deleted if a user is restoring or retrieving a file residing in the volume, if the server is writing information to the volume, or if a reclamation process is using the volume.

If you issue the DELETE VOLUME command, volume information is deleted from the IBM Spectrum Protect™ database. However, the physical files that are allocated with DEFINE VOLUME command are not removed from the file space.

If this command is applied to a WORM (write once, read many) volume, the volume returns to scratch if it has space remaining in which data can be written. Data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can only be written in space that does not contain current, deleted, or expired data. If a WORM volume does not have any space available in which data can be written, it remains private. To remove the volume from the library, you must use the CHECKOUT LIBVOLUME command.

The DELETE VOLUME command automatically updates the server library inventory for sequential volumes if the volume is returned to scratch status when the volume becomes empty. To determine whether a volume will be returned to scratch status, issue the QUERY VOLUME command and look at the output. If the value for the attribute "Scratch Volume?" is "Yes," then the server library inventory is automatically updated.

If the value is "No," you can issue the UPDATE LIBVOLUME command to specify the status as scratch. It is recommended that you issue the UPDATE LIBVOLUME command after issuing the DELETE VOLUME command.

Attempting to use the DELETE VOLUME command to delete WORM FILE volumes in a storage pool with RECLAMATIONTYPE=SNAPLOCK fails with an error message. Deletion of empty WORM FILE volumes is performed only by the reclamation process.

If you issue the DELETE VOLUME command for a volume in a storage pool that has a SHRED parameter value greater than 0, the volume is placed in the pending state until shredding is run. Shredding is necessary to complete the deletion, even if the volume is empty.

If you issue the DELETE VOLUME command for a volume in a storage pool that is set up for data deduplication, the server destroys any object that is referencing data on that volume.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax

```
.-DISCARDdata-----No-----.
```

```

>>-DELEte Volume--volume_name--+-----+----->
                                     '-DISCARDdata----+No--+-'
                                     '-Yes-'

.-Wait-----No-----.
>--+-----+----->>
   '-Wait-----+No--+-'
   '-Yes-'

```

Parameters

volume_name (Required)

Specifies the name of the volume to delete.

DISCARDdata

Specifies whether files stored in the volume are deleted. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that files stored in the volume are not deleted. If the volume contains any files, the volume is not deleted.

Yes

Specifies that all files stored in the volume are deleted. The server does not need to mount the volume for this type of deletion.

Remember:

1. The server does not delete archive files that are on deletion hold.
2. If archive retention protection is enabled, the server deletes only archive files whose retention period has expired.

If the volume being deleted is a primary storage pool volume, the server checks whether any copy storage pool has copies of files that are being deleted. When files stored in a primary storage pool volume are deleted, any copies of these files in copy storage pools are also deleted.

When you delete a disk volume in a primary storage pool, the command also deletes any files that are cached copies (copies of files that have been migrated to the next storage pool). Deleting cached copies of files does not delete the files that have already been migrated or backed up to copy storage pools. Only the cached copies of the files are affected.

If the volume being deleted is a copy storage pool volume, only files on the copy pool volume are deleted. The primary storage pool files are not affected.

Do not use the DELETE VOLUME command with DISCARDDATA=YES if a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The DELETE VOLUME command could cause the restore to be incomplete.

If you cancel the DELETE VOLUME operation during processing or if a system failure occurs, some files might remain on the volume. You can delete the same volume again to have the server delete the remaining files and then the volume.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter affects processing only when you have also requested that any data on the volume be discarded. This parameter is optional. The default value is No. Possible values are:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Remember: You cannot specify WAIT=YES from the server console.

Example: Delete a storage pool volume

Delete storage pool volume stgvol.1 from the storage pool FILEPOOL.

```
delete volume stgvol.1
```

Related commands

Table 1. Commands related to DELETE VOLUME

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

DISABLE commands

Use DISABLE commands to prevent some types of operations by the server.

- DISABLE EVENTS (Disable events for event logging)
- DISABLE REPLICATION (Prevent outbound replication processing on a server)
- DISABLE SESSIONS (Prevent new sessions from accessing IBM Spectrum Protect)

DISABLE EVENTS (Disable events for event logging)

Use this command to disable the processing of one or more events. If you specify a receiver that is not supported on any platform, or if you specify an invalid event or name, IBM Spectrum Protect™ issues an error message. However, any valid receivers, events, or names that you specified are still enabled.

Tip: Messages in the SEVERE category and message ANR9999D can provide valuable diagnostic information if there are serious server problems. For this reason, you should not disable these messages.

Restriction:

- Certain messages are displayed on the console even if they are disabled. These include some messages issued during server startup and shutdown and responses to administrative commands.
- Server messages from the server on which this command is issued cannot be disabled for the activity log.

ANR1822I indicates that event logging is being ended for the specified receiver. When the DISABLE EVENTS command is issued, this message is logged to the receiver even if it is one of the events that has been disabled. This is done to confirm that event logging has ended to that receiver, but subsequent ANR1822I messages are not logged to that receiver.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
.-.,----- .-.,-----
```

```

      v                | v                |
>>-DISable EVents-----+receivers-----+-----+event_name+----->
      +-ALL-----+          +-ALL-----+
      +-CONSOLE-----+        +-INFO-----+
      +-ACTLOG-----+        +-WARNING-----+
      +-EVENTSERVER-----+    +-ERROR-----+
      +-FILE-----+          '-SEVERE-----'
      +-FILETEXT-----+
      |                (1) |
      +-NTEVENTLOG-----+
      |                (2) |
      +-SYSLOG-----+
      +-TIVOLI-----+
      '-USEREXIT-----'

>----->
|                .-,----- .                |
|                v                |
+-NODEname-----+node_name+-----+
|                .-,----- .                |
|                v                |
'-SERVername-----+server_name+-----'

```

Notes:

1. NTEVENTLOG is available only on Windows.
2. SYSLOG is available only on Linux.

Parameters

receivers (Required)

Specifies the name of the receivers for which to disable events. Specify multiple receivers by separating them with commas and no intervening spaces. Possible values are:

ALL

All receivers, except for server events on the activity log receiver (ACTLOG). Only client events can be disabled for the activity log receiver.

CONSOLE

The standard server console as a receiver.

ACTLOG

The activity log as a receiver. You can disable only client events, not server events, for the activity log.

EVENTSERVER

The event server as a receiver.

FILE

A user file as a receiver. Each logged event is a record in the file. The records are not easily readable by people.

FILETEXT

A user file as a receiver. Each logged event is a fixed-size, readable line.

NTEVENTLOG

The Windows application log as a receiver.

Linux **SYSLOG**

Linux Writes messages directly to the system log on Linux.

TIVOLI

The Tivoli Enterprise Console® (TEC) as a receiver.

USEREXIT

A user-written program as a receiver. The server writes information to the program.

events (Required)

Specifies the events to be disabled. You can specify multiple events by separating them with commas and no intervening spaces. Possible values are:

ALL

All events.

event_name

A four-digit message number preceded by **ANR** for a server event or **ANE** for a client event. Valid ranges are from ANR0001 to ANR9999 and from ANE4000 to ANE4999. Specify the NODENAMES parameter if client events are to

be disabled for matching nodes. Specify the SERVERNAME parameter if server events are to be disabled for matching servers.

For the TIVOLI event receiver only, you can specify the following events names for the IBM Spectrum Protect application clients:

IBM Spectrum Protect application client	Prefix	Range
Data Protection for Microsoft Exchange Server	ACN	3500–3649
Data Protection for Lotus® Domino®	ACD	5200–5299
Data Protection for Oracle	ANS	500–599
Data Protection for Informix®	ANS	600–699
Data Protection for Microsoft SQL Server	ACO	3000–3999

Remember: Specifying ALL disables these messages. However, the INFO, WARNING, ERROR, and SEVERE options have no effect on the messages.

severity categories

If the event list contains a severity category, all events of that severity are disabled for the specified nodes. The message types are:

INFO

Information messages (type of I).

WARNING

Warning messages (type of W).

ERROR

Error messages (type of E).

SEVERE

Severe error messages (type of S).

NODENAME

Specifies the name of one or more node names for which events are to be disabled. You can use the wildcard character (*) to specify all nodes. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the events are disabled for the server running this command.

SERVername

Specifies the name of one or more server names for which events are to be disabled. You can use the wildcard character (*) to specify all servers other than the server running this command. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the events are disabled for the server running this command.

Example: Disable specific categories of events

Disable all client events in the INFO and WARNING categories for the activity log and console receivers for all nodes.

```
disable events actlog,console
info,warning nodename=*
```

Related commands

Table 1. Commands related to DISABLE EVENTS

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

DISABLE REPLICATION (Prevent outbound replication processing on a server)

Use this command to prevent a source replication server from starting new replication processes.

The use of this command does not stop running replication processes. Running replication processes continue until they complete or until they end without completing. Use this command and the ENABLE REPLICATION command to control replication processing.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-DISAbLe REPLIcation-----<<
```

Parameters

None.

Example: Disable replication processing

Disable replication processing on a source replication server.

```
disable replication
```

Related commands

Table 1. Commands related to DISABLE REPLICATION

Command	Description
CANCEL REPLICATION	Cancels node replication processes.
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
ENABLE REPLICATION	Allows outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

DISABLE SESSIONS (Prevent new sessions from accessing IBM Spectrum Protect)

Use this command to prevent new sessions from accessing IBM Spectrum Protect™. Active sessions will complete. For a particular server, you can specify whether to disable inbound sessions, outbound sessions, or both.

Server processes, such as migration and reclamation, are not affected when you issue the DISABLE SESSIONS command.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
>>-DISAbLe SESSions----->
.-CLient-----
>--+-----+><
'|+-CLient-----+'
  +-ALL-----+
  +-ADMin-----+
  '-SERVer-----+'
    | .-DIRectioN---Both-----|
    '-server_name-----+'
      '+-DIRectioN---Both-----+'
      +-DIRectioN---INbound---+
      '-DIRectioN---OUTbound-'
```

Parameters

Specifies the type of session to be disabled. This parameter is optional. The default value is CLIENT. You can specify one of the following values:

CLient

Disables only backup and archive client sessions.

ALL

Disables all session types.

ADMin

Disables only administrative sessions.

SERVer

Disables only server-to-server sessions. Only the following types of sessions are disabled:

- Server-to-server event logging
- Enterprise management
- Server registration
- LAN-free: storage agent - server
- Virtual volumes
- Node replication

You can also specify whether to disable inbound sessions, outbound sessions, or both for a particular server.

server_name

Specifies the name of a server whose sessions you want to disable. This parameter is optional. If you do not specify this parameter, new sessions with other servers do not start. Running sessions are not canceled.

DIRectioN

Specifies whether to disable inbound sessions, outbound sessions, or both. This parameter is optional. The default is BOTH. The following values are possible:

Both

Specifies that inbound sessions from the specified server and outbound sessions to the specified server are disabled.

INbound

Specifies that only inbound sessions from the specified server are disabled.

OUTbound

Specifies that only outbound sessions to the specified server are disabled.

Example: Prevent new client node backup and archive sessions on the server

Temporarily prevent new client node sessions from accessing the server.

```
disable sessions
```

Example: Prevent all new sessions on the server

Temporarily prevent any new sessions from accessing the server.

```
disable sessions all
```

Example: Disable outbound sessions to a server

Disable outbound sessions to a server named REPLSRV.

```
disable sessions server replsrv direction=outbound
```

Related commands

Table 1. Commands related to DISABLE SESSIONS

Command	Description
CANCEL SESSION	Cancels active sessions with the server.
DISABLE REPLICATION	Prevents outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

DISMOUNT command

Use the DISMOUNT command to dismount a volume by the real device address or by volume name.

- DISMOUNT VOLUME (Dismount a volume by volume name)

DISPLAY OBJNAME (Display a full object name)

Use this command when you want IBM Spectrum Protect™ to display a full object name if the name displayed in a message or query output has been abbreviated due to length. Object names that are very long can be difficult to display and use through normal operating system facilities. The IBM Spectrum Protect server will abbreviate long names and assign them a token ID which might be used if the object path name exceeds 1024 bytes. The token ID is displayed in a string that includes identifiers for the node, file space, and object name. The format is: [TSMOBJ:*nID.fsID.objID*]. When specified with the DISPLAY OBJNAME command, the token ID can be used to show the full object name.

Privilege class

Any administrator can issue this command

Syntax

```
>>-DISplay OBJname--token_ID-----<<
```

Parameters

token_ID (Required)

Specifies the ID reported in the [TSMOBJ:] tag, when an object name is too long to display.

Example: Display the full object name of a token ID in a message

Assume the you receive the following message:

```
ANR9999D file.c(1999) Error handling file [TSMOBJ:1.1.649498] because  
of lack of server resources.
```

Display the full object name for the file referenced in the error message by specifying the token ID on the DISPLAY OBJNAME command.

```
display obj 1.1.649498
```

Related commands

Table 1. Commands related to DISPLAY OBJNAME

Command	Description
QUERY CONTENT	Displays information about files in a storage pool volume.

ENABLE commands

Use ENABLE commands to allow some types of operations by the server.

- ENABLE EVENTS (Enable server or client events for logging)
- ENABLE REPLICATION (Allow outbound replication processing on a server)
- ENABLE SESSIONS (Resume user activity on the server)

ENABLE EVENTS (Enable server or client events for logging)

Use this command to enable the processing of one or more events. If you specify a receiver that is not supported on any platform, or if you specify an invalid event or name, IBM Spectrum Protect™ issues an error message. However, any valid receivers, events, or names that you specified are still enabled.

Restriction: Certain events, such as some messages issued during server start-up and shutdown, automatically go to the console. They do not go to other receivers even if they are enabled.

Administrative commands are returned to the command issuer and are only logged as numbered events. These numbered events are not logged to the system console, but are logged to other receivers, including administrative command-line sessions running in console mode.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-,----- .-,-----
      v          |          v          |
>>-ENable--EEvents---+--ALL-----+-----+event_name+---->
      +-CONSOLE-----+          +-ALL-----+
      +-ACTLOG-----+          +-INFO-----+
      +-EVENTSERVER----+          +-WARNING----+
      +-FILE-----+          +-ERROR-----+
      +-FILETEXT-----+          '-SEVERE-----'
      |                (1) |
      +-NTEVENTLOG-----+
      |                (2) |
      +-SYSLOG-----+
      +-TIVOLI-----+
      '-USEREXIT-----'

>--|-----|-----<<
|          .-,----- .-,-----
|          v          |          v          |
+-NODEname-----+node_name+-----+
|          .-,----- .-,-----
|          v          |          v          |
|'-SERVername-----+server_name+--'

```

Notes:

1. NTEVENTLOG is available only on Windows.
2. This parameter is only available for the Linux operating system.

Parameters

receivers (Required)

Specifies one or more receivers for which to log enabled events. You can specify multiple receivers by separating them with commas and no intervening spaces. Valid values are:

ALL

All receivers.

CONSOLE

The standard server console as a receiver.

ACTLOG

The server activity log as a receiver.

EVENTSERVER

The event server as a receiver.

FILE

A user file as a receiver. Each logged event is a record in the file. The records are not easily readable by people.

FILETEXT

A user file as a receiver. Each logged event is a fixed-size, readable line.

Windows NTEVENTLOG

The Windows application log as a receiver.

Linux SYSLOG

Specifies the Linux system log as a receiver with a facility of LOG_USER.

TIVOLI

The Tivoli Enterprise Console® (TEC) as a receiver.

USEREXIT

A user-written program as a receiver. The server writes information to the program.

events (Required)

Specifies the type of events to be enabled. You can specify multiple events by separating them with commas and no intervening spaces. Possible values are:

ALL

All events.

event_name

A four-digit message number preceded by ANR for a server event or ANE for a client event. Valid ranges are from ANR0001 to ANR9999 and from ANE4000 to ANE4999. Specify the NODENAME parameter if client events are to be enabled for matching nodes. Specify the SERVERNAME parameter if server events are to be enabled for matching servers.

For the TIVOLI event receiver, you can specify the following additional ranges for the IBM Spectrum Protect application clients:

IBM Spectrum Protect application client	Prefix	Range
Data Protection for Microsoft Exchange Server	ACN	3500–3649
Data Protection for Lotus® Domino®	ACD	5200–5299
Data Protection for Oracle	ANS	500–599
Data Protection for Informix®	ANS	600–699
Data Protection for Microsoft SQL Server	ACO	3000–3999

Restriction: The application client must have enhanced Tivoli® Event Console support enabled in order to route these messages to the Tivoli Event Console.

Tip:

- Specifying the ALL option enables these messages. However, the INFO, WARNING, ERROR, and SEVERE options have no effect on the messages.
- Because of the number of messages, you should not enable all messages from a node to be logged to the Tivoli Event Console.

severity categories

If the event list contains a severity category, all events of that severity are enabled for the specified nodes. The message types are:

INFO

Information messages (type of I) are enabled.

WARNING

Warning messages (type of W) are enabled.

ERROR

Error messages (type of E) are enabled.

SEVERE

Severe error messages (type of S) are enabled.

NODENAME

Specifies one or more client nodes for which events are enabled. You can use a wildcard character to specify all client nodes. You can specify NODENAME or SERVERNAME. If neither parameter is specified, events are enabled for the server running this command.

SERVERNAME

Specifies one or more servers for which events are to be enabled. You can use a wildcard character to specify all servers other than the server from which this command is issued. You can specify SERVERNAME or NODENAME. If neither parameter is specified, the events are enabled for the server running this command.

Example: Enable specific categories of events

Enable all ERROR and SEVERE client events to the USEREXIT receiver for the node BONZO.

```
enable events userexit error,severe nodename=bonzo
```

Related commands

Table 1. Commands related to ENABLE EVENTS

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

ENABLE REPLICATION (Allow outbound replication processing on a server)

Use this command to allow a source replication server to begin normal replication processing after a database restore. You can also use this command to resume replication processing after issuing the DISABLE REPLICATION command.

Attention: Before enabling replication after a database restore, determine whether copies of data that are on the target server are needed. If they are, you must synchronize client node data by replicating the data from the target replication server to the source replication server. The replication process replaces the data on the source server that was lost because of the database restore.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-ENable REPLication----->>
```

Parameters

None.

Example: Allow replication processing

Allow replication processing on a source replication server.

```
enable replication
```

Related commands

Table 1. Commands related to ENABLE REPLICATION

Command	Description
DISABLE REPLICATION	Prevents outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

ENABLE SESSIONS (Resume user activity on the server)

Use this command after issuing the DISABLE SESSIONS command to start new sessions that can access a server. For a particular server, you can specify whether to enable inbound sessions, outbound sessions, or both.

The processing of this command does not affect system processes, such as migration and reclamation.

Use the QUERY STATUS command to display the availability of the server.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
>>-ENable SESSions----->>
.-CLient-----
>--+-----+>>
'|+-CLient-----+'
'|+-ALL-----+'
'|+-ADMin-----+'
'| -SERVer--+-----+'
'|          |          .-DIRection----Both-----.'
'| -server_name--+-----+'
'|          |          +-DIRection----Both-----+'
'|          |          +-DIRection----INbound--+
'|          |          '-DIRection----OUTbound-'
```

Parameters

Specifies the type of session to be enabled. This parameter is optional. The default value is CLIENT. You can specify one of the following values:

CLient

Enables only backup and archive client sessions.

ALL

Enables all session types.

ADMin

Enables only administrative sessions.

SERVer

Enables only server-to-server sessions. You can also specify whether to enable inbound sessions, outbound sessions, or both for a particular server.

server_name

Specifies the name of a particular server whose sessions you want to enable. This parameter is optional. If you do not specify this parameter, new sessions with all other servers are enabled.

DIRection

Specifies whether to enable inbound sessions, outbound sessions, or both. This parameter is optional. The default is BOTH. The following values are possible:

Both

Specifies that inbound sessions from the specified server and outbound sessions to the specified server are enabled.

INbound

Specifies that only inbound sessions to the specified server are enabled.

OUTbound

Specifies that only outbound sessions from the specified server are enabled.

Example: Resume client node activity on the server

Resume normal operation, permitting client nodes to access the server.

```
enable sessions
```

Example: Resume all activity on the server

Resume normal operation, permitting all sessions to access the server.

```
enable sessions all
```

Example: Enable outbound sessions to a server

Enable outbound sessions to a server named REPLSRV.

```
enable sessions server replsrv direction=outbound
```

Related commands

Table 1. Commands related to ENABLE SESSIONS

Command	Description
ACCEPT DATE	Accepts the current date on the server.
CANCEL SESSION	Cancels active sessions with the server.
ENABLE REPLICATION	Allows outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

ENCRYPT STGPOOL (Encrypt data in a storage pool)

Use this command to encrypt data in a directory-container or cloud-container storage pool.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-ENcRypt STGpooL--pool_name--+-----+----->
                                     .-MAXPRocess---4-----
                                     '-MAXPRocess---number-'

.-Preview---No----- .-Wait---No-----
>--+-----+----->>
'-Preview---+Yes-+-' '-Wait---+No-+-'
      '-No--'          '-Yes-'
```

Parameters

pool_name (Required)

Specifies the name of the storage pool that contains data that must be encrypted.

Restrictions:

- You can specify only directory-container storage pools or cloud-container storage pools.
- You can specify up to 30 characters for the storage pool name. If you specify more than 30 characters, the command fails.

MAXPRocess

Specifies the maximum number of parallel processes that can occur when the storage pool is encrypting data. This parameter is optional. Enter a value in the range 1 - 99. The default value is 4.

Preview

Specifies whether a preview is displayed of all the commands that are processed as part of the ENCRYPT STGPOOL command. This parameter is optional. The following values are possible:

No

Specifies that a preview of the commands is not displayed. This is the default value.

Yes

Specifies that a preview of the commands is displayed.

Wait

Specifies whether the storage pool encryption occurs in the foreground or background. This parameter is optional. You can specify one of the following values:

No

Specifies that the operation is completed in the background. You can continue with other tasks while the command is processing. Messages that are related to the background process are displayed in the activity log file or the server console, depending on where the messages are logged. This is the default value.

Yes

Specifies that the operation is completed in the foreground. It might take a long time to complete the operation. The operation must end before you can continue with other tasks. Messages are displayed in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify the WAIT=YES parameter from the server console.

Example: Encrypt data in a storage pool

Encrypt data in a storage pool that is named POOL1 and specify a maximum number of 30 parallel processes.

```
encrypt stgpool pool1 maxprocess=30
```

Related commands

Table 1. Commands related to ENCRYPT STGPOOL

Command	Description
DEFINE STGPOOL (directory-container)	Define a directory-container storage pool.

END EVENTLOGGING (Stop logging events)

Use this command to stop logging events to an active receiver.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-END--Eventlogging--+-ALL-----+----->>
| .-,-----+-----+----->>
| V-----+-----+----->>
'---+-CONSOLE-----+---'
+-ACTLOG-----+
+-EVENTSERVER----+
+-FILE-----+
+-FILETEXT-----+
| (1) |
+-NTEVENTLOG-----+
| (2) |
+-SYSLOG-----+
+-TIVOLI-----+
'-USEREXIT-----'

```

Notes:

1. This parameter is only available for Windows operating system.
2. This parameter is only available for the Linux operating system.

Parameters

Specify a type of receiver. You can specify multiple receivers by separating them with commas and no intervening spaces. This is an optional parameter. The default is ALL. If you specify ALL or no receiver, logging ends for all receivers.

ALL

Specifies all receivers.

CONSOLE

Specifies the server console as a receiver.

ACTLOG

Specifies the IBM Spectrum Protect™ activity log as a receiver. Logging can be stopped only for client events.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

Windows NTEVENTLOG

Windows Specifies the Windows application log as a receiver.

Linux SYSLOG

Linux Specifies the Linux system log as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which IBM Spectrum Protect writes information as a receiver.

Example: Stop logging events

End logging of events to the user exit.

```
end eventlogging userexit
```

Related commands

Table 1. Commands related to END EVENTLOGGING

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

EXPIRE INVENTORY (Manually start inventory expiration processing)

Use this command to manually start inventory expiration processing. The inventory expiration process removes client backup and archive file copies from server storage. Removal is based on policy specifications in the backup and archive copy groups of the management classes to which the files are bound.

When you have the disaster recovery manager function for your IBM Spectrum Protect™ server, the inventory expiration process also removes eligible virtual volumes that are used by the following processes:

- Database backups of type BACKUPFULL, BACKUPINCR, and DBSNAPSHOT. The SET DRMDBBACKUPEXPIREDDAYS command controls when these volumes are eligible for expiration.
- Recovery plan files of type RPFIL and RPFNSNAPSHOT. The SET DRMRPFEXPIREDDAYS command controls when these volumes are eligible for expiration.

The inventory expiration process that runs during server initialization does not remove these virtual volumes.

Only one expiration process is allowed at any time, but this process can be distributed among a maximum of 40 threads. If an expiration process is running, you cannot start another process.

You can set up automatic expiration processing with the EXPINTERVAL server option. If you set the EXPINTERVAL option to 0, the server does not run expiration automatically, and you must issue the EXPIRE INVENTORY command to start expiration processing.

This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information about background processes, use the QUERY PROCESS command.

If this command is applied to a WORM volume, the volume returns to being a scratch volume if it has remaining space in which data can be written. Data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can be written only in space that does not contain current, deleted, or expired data. If a WORM volume does not have any space available in which data can be written, it remains private. To remove the volume from the library, you must use the CHECKOUT LIBVOLUME command.

Run the EXPIRE INVENTORY command to delete files from server storage if they were not deleted when you used client delete operations.

For more information about client delete operations, see Backup-archive client options and commands.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

.-Quiet-----No-----.
>>-EXPIre Inventory----->
'-Quiet-----+No--+'
'-Yes-'

.-Wait-----No-----.  .-Nodes-----*------.
>----->
'-Wait-----+No--+'  '-Nodes-----+node_name-----+'
'-Yes-'                '-node_group_name-'

>----->
'-EXCLUDENodes-----excluded_node_name-'

.-Type-----All-----.
>----->
'-Domain-----domain_name-'  '-Type-----+All-----+'
                               +-Archive-+
                               +-Backup--+
                               '-Other---'

.-Resource-----4-----.  .-Skipdirs-----No-----.
>----->
'-Resource-----number-'  '-Skipdirs-----+No--+'
                               '-Yes-'

>----->>
'-Duration-----minutes-'

```

Parameters

Quiet

Specifies whether the server suppresses detailed messages about policy changes during the expiration processing. This parameter is optional. The default is NO. Possible values are:

No

Specifies that the server sends detailed informational messages.

Yes

Specifies that the server sends only summary messages. The server issues messages about policy changes only when files are deleted and either the default management class or retention grace period for the domain was used to expire the files.

You can also specify the EXPQUIET option in the server options file to automatically determine whether expiration processing is run with summary messages.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before you continue with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

Skipdirs

Specifies whether the server skips directory type objects during the expiration processing. This parameter is optional. The default is NO. Possible values are:

No

Specifies that the server expires files and directories that are based on the appropriate policy criteria.

Yes

Specifies that the server skips directory type backup and archive objects during expiration processing, even if the directories are eligible for expiration. By specifying YES, you prevent deletion of directories, and expiration processing can occur more quickly.

Attention: Do not use this option all of the time. With IBM Spectrum Protect Version 6.0 and later, you can run multiple threads (resources) for an expiration process. Also, if you specify YES often, the database grows as the directory objects accumulate, and the time that is spent for expiration increases. Run SKIPDIRS=NO periodically to expire the directories and reduce the size of the database.

Nodes

Specifies the name of the client nodes or node groups whose data is to be processed. To specify multiple node and node group names, separate the names with commas and no intervening spaces. Node names can contain wildcard characters, but node group names cannot. This parameter is optional.

You can specify NODES, EXCLUDENODES, DOMAIN, or any combination. If you specify more than one of these parameters, only those nodes that match the criteria for both NODES and DOMAIN and does not match the criteria for EXCLUDENODES command options are processed. If you do not specify NODES, EXCLUDENODES, or DOMAIN with a value, data for all nodes is processed.

EXCLUDENodes

Specifies the name of the client nodes or node groups whose data is not to be processed. To specify multiple node and node group names, separate the names with commas and no intervening spaces. Node names can contain wildcard characters, but node group names cannot. This parameter is optional.

You can specify NODES, EXCLUDENODES, DOMAIN, or any combination. If you specify more than one of these parameters, only those nodes that match the criteria for both NODES and DOMAIN and does not match the criteria for EXCLUDENODES command options are processed. If you do not specify NODES, EXCLUDENODES, or DOMAIN with a value, data for all nodes is processed.

Domain

Specifies that only data for client nodes that are assigned to the specified domain is to be processed. This parameter is optional. You can specify NODES, EXCLUDENODES, DOMAIN, or any combination. If you specify more than one of these parameters, only those nodes that match the criteria for both NODES and DOMAIN and does not match the criteria for EXCLUDENODES command options are processed. If you do not specify NODES, EXCLUDENODES, or DOMAIN with a value, data for all nodes is processed.

Type

Specifies the type of data to be processed. This parameter is optional. The default value is ALL. Possible values are:

ALL

Process all types of data that is eligible for expiration

Archive

Process only client archive data

Backup

Process only client backup data

Other

Process only items for disaster recovery manager functions, such as recovery plan files and obsolete database backups

REsource

Specifies the number of threads that can run in parallel. Specify a value in the range 1 - 40. This parameter is optional. The default is four.

Expiration runs as a single process, although the resources represent parallel work by the server within the single expiration process. Archive data for a node runs only on a single resource, but backup data can be spread across resources on a file space level. For example, if you specify NODE=X, Y, Z each with three file spaces and RESOURCE=5, then expiration processing for the three X, Y, and Z client nodes runs in parallel. At least one resource processes each node, and at least one node uses multiple resources for processing backup data across the multiple file spaces.

DURATION

Specifies the maximum number of minutes for the expiration process to run. The process stops when the specified number of minutes pass or when all eligible expired objects are deleted, whichever comes first. Specify a value in the range 1 - 2880. This parameter is optional. If this parameter is not specified, the duration of the expiration process is not limited by time.

Example: Run inventory expiration processing for a specific time period

Run the expiration process for two hours.

```
expire inventory duration=120
```

Example: Run inventory expiration processing for backup data for two client nodes

Run inventory expiration processing for the backup data for two client nodes, CHARLIE and ROBBIE. Allow the server to run expiration processing until completed.

```
expire inventory nodes=charlie,robbie resource=2 type=backup
```

Example: Run inventory expiration processing for all client nodes except two nodes

Run inventory expiration processing for all client nodes except two nodes, CHARLIE and ROBBIE. Allow the server to run expiration processing until completed.

```
expire inventory excludenodes=charlie,robbie
```

Example: Run inventory expiration processing for all client nodes in a domain except one node

Run inventory expiration processing for all client nodes in a domain except one node, ROBBIE. Allow the server to run expiration processing until completed.

```
expire inventory domain=standard excludenodes=robbie
```

Related commands

Table 1. Commands related to EXPIRE INVENTORY

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
CANCEL EXPIRATION	Cancels inventory expiration processing.
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background processes.

EXPORT commands

Use the EXPORT commands to copy information from an IBM Spectrum Protect™ server to sequential removable media.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the EXPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

- EXPORT ADMIN (Export administrator information)
- EXPORT NODE (Export client node information)
- EXPORT POLICY (Export policy information)
- EXPORT SERVER (Export server information)

EXPORT ADMIN (Export administrator information)

Use this command to export administrator and authority definitions from a server. You can export the information to sequential media for later importing to another server, or you can export the information directly to another server for immediate import.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect™ server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already

synchronized by that server, you must update the password. After issuing the EXPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

IBM Spectrum Protect exports administrator information such as:

- Administrator name, password, and contact information
- Administrative privilege classes that are granted to the administrator
- Whether the administrator ID is locked from server access

You can use the QUERY ACTLOG command to view the status of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If you export information to sequential media and the background process is canceled, the sequential media that is holding the exported data is incomplete, it must not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, use the QUERY PROCESS command.

The following restrictions apply to the export function:

- Export operations from a later version and release to an earlier version and release is not supported.
- Export operations between servers that are at the same version and release but with different fix packs might fail. For example, you cannot export from a V7.1.3 server to a V7.1.1 or earlier server.
- Exported data from a server with retention protection enabled is not protected by retention when it is imported to another server.
- Export processing excludes nodes of type network-attached storage (NAS).
- Exporting data to a Centera device class or importing data from a Centera device class is not supported. However, files that are stored in Centera storage pools can be exported and files that must be imported can be stored on a Centera storage device.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The EXPORT ADMIN command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

Table 1. Commands related to EXPORT ADMIN

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

- EXPORT ADMIN (Export administrator definitions to sequential media)
You can export administrator and authority definitions from a server to sequential media for later importing to another server.
- EXPORT ADMIN (Export administrator information directly to another server)
Use this command to export administrator and authority definitions directly to another server on the network. This results in an immediate import on the target server.

EXPORT ADMIN (Export administrator definitions to sequential media)

You can export administrator and authority definitions from a server to sequential media for later importing to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-EXPort Admin-----*----->
| .,-----|
| V        |
|'---admin_name---'|

.-Preview---No----->
| (1) (2) |
|'Preview-----No---'|
|          '-Yes- '|

>----->
| (1) |
|'DEVclass-----device_class_name-'|

.-Scratch---Yes----->
| (2) |
|'Scratch-----Yes---'|
|          '-No-- '|

>----->
| .,-----|
| (2) V        |
|'VOLumentnames-----+---volume_name+---'|
|          '-FILE:--file_name-'|

>----->
|'USEDVolumelist-----file_name-'|

.-ENCryptionstrength---AES----->
|'ENCryptionstrength-----+AES---'|
|          '-DES- '|
```

Notes:

1. If PREVIEW=NO, a device class must be specified.
2. If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

Parameters

admin_name

Specifies the administrators for which information is to be exported. This parameter is optional. The default is all administrators.

Separate the items in the list by commas, with no intervening spaces. You can use wildcard characters to specify names.

Preview

Specifies whether to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data are transferred, and determine how many volumes are required. The following parameter values are supported:

No

Specifies that the administrator information is to be exported. If you specify this value, you must specify a device class.

Yes

Specifies that the operation is previewed but not completed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

This parameter is optional. The default value is NO.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the export runs, IBM Spectrum Protect™ cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER.

Scratch

Specifies whether scratch volumes can be used. The default value is YES. You can specify one of the following values:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No

Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumenames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

You can specify one of the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1-6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: <div style="display: flex; justify-content: space-between; margin-top: 5px;"> AIX Linux /imdata/mt1. </div> <div style="display: flex; justify-content: space-between; margin-top: 5px;"> Windows d:\program files\tivoli\tsm\data1.dsm. </div>
<div style="display: flex; justify-content: space-between; margin-bottom: 5px;"> AIX Linux Windows </div> REMOVABLEFILE	<div style="display: flex; justify-content: space-between; margin-bottom: 5px;"> AIX Linux Windows </div> 1-6 alphanumeric characters.
SERVER	1-250 alphanumeric characters.

USEDVolumelist

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. You can specify one of the following values:

AES

- Specifies the Advanced Encryption Standard.
- DES
Specifies the Data Encryption Standard.

Example: Export administrator definitions to tape volumes

From the server, export the information for all defined administrators to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class. The number and types of objects that are exported are reported to the system console and in the activity log. Issue the command:

```
export admin devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Export administrator definitions to tape volumes listed in a file

From the server, export the information for all defined administrators to tape volumes that are listed in the following file:

- **AIX** | **Linux** TAPEVOL
- **Windows** TAPEVOL.DATA

This file contains the following lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that these tape volumes be used by a device that is assigned to the MENU1 device class. Issue the command:

```
AIX | Linux
export admin devclass=menu1 volumenames=file:tapevol

Windows
export admin devclass=menu1 volumenames=file:tapevol.data
```

The number and types of objects that are exported are reported to the system console and in the activity log.

EXPORT ADMIN (Export administrator information directly to another server)

Use this command to export administrator and authority definitions directly to another server on the network. This results in an immediate import on the target server.

You can issue a QUERY PROCESS command from the target server to monitor the progress of the import operation. See EXPORT ADMIN (Export administrator information) for a list of restrictions that apply to the export function.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-EXPort Admin-----*-----
| .-,-----|
| V          |
|'---admin_name-+-'

>-----PREVIEWImport-----No-----
|'---TOserver-----servername-' |'---PREVIEWImport-----+No---+'
|                                     |'---Yes-'

>-----Replacedefs-----No-----
|'---Replacedefs-----+No---+'

```


EXPORT NODE (Export client node information)

Use this command to export client node definitions or file data to sequential media or directly to another server for immediate import.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect™ server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the EXPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

The following information is included in each client node definition:

- User ID, password, and contact information.
- Name of the client's assigned policy domain.
- File compression status.
- Whether the user has the authority to delete backed-up or archived files from server storage.
- Whether the client node ID is locked from server access.

Optionally, you can also export the following items:

- File space definitions.
- Backed-up, archived, and files that were migrated by an IBM Spectrum Protect for Space Management client.
- Access authorization information that pertains to the file spaces exported.
- Archive data that is in deletion hold status (the hold status is preserved). When the archive data is imported, it remains in deletion hold.

If you use an LDAP directory server to authenticate passwords, any servers that you export to must be configured for LDAP passwords. Node data that is exported from a node that authenticates with an LDAP directory server is inaccessible if the target server is not properly configured. If your target server is not configured, exported data from an LDAP node can still be exported. But the target server must be configured to use LDAP, to access the data.

The following restrictions apply to the export function:

- Export operations from a later version and release to an earlier version and release is not supported.
- Export operations between servers that are at the same version and release but with different fix packs might fail. For example, you cannot export from a V7.1.3 server to a V7.1.1 or earlier server.
- Exported data from a server with retention protection enabled is not protected by retention when it is imported to another server.
- Export processing excludes nodes of type network-attached storage (NAS).
- Exporting data to a Centera device class or importing data from a Centera device class is not supported. However, files that are stored in Centera storage pools can be exported and files that must be imported can be stored on a Centera storage device.
- The EXPORT NODE and EXPORT SERVER commands do not export data from a shred pool unless you explicitly allow it by setting the ALLOWSHREDDABLE parameter to the YES value. If this value is specified, and the exported data includes data from shred pools, that data cannot be shredded. A warning is not issued if the export operation includes data from shred pools.
- Incrementally exporting or importing the following types of client data to another IBM Spectrum Protect server is not supported:
 - VMware backups where full plus incremental backups need to be periodically, incrementally transferred to another server
 - Backups groups where full plus differential backups must be periodically, incrementally transferred to another server
 - Windows System State data that is periodically, incrementally transferred to another server

Full export or import of this data to a new file system on the target is supported by exporting the entire file space that contains the data. The export must not use the FILEDATA=ALLACTIVE, FROMDATE, TODATE, or MERGEFILESPPACES parameters.

Using node replication to incrementally transfer this type of client data between two servers is optimal.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The EXPORT NODE command generates a background process that can be canceled with the CANCEL PROCESS command. If you are exporting node information to sequential media and the background process is canceled, the sequential media that is holding the exported data is incomplete, it must not be used to import data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, issue the QUERY PROCESS command.

To display information about any running and suspended server-to-server export operations, issue the QUERY EXPORT command. The QUERY EXPORT command displays information only for exports that are, or can be, suspended. Export operations that can be suspended, and then restarted, are those server-to-server exports whose FILEDATA has a value other than NONE. You can issue the QUERY ACTLOG command to view the status of the export operation.

Because of unpredictable results, do not run expiration, migration, backup, or archive when you are issuing the EXPORT NODE command.

For a server that has clients with support for Unicode, you can get the server to convert the file space name that you enter, or use one of the following parameters:

- FSID
- UNIFILESPACE

The EXPORT NODE command takes two forms: export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

Table 1. Commands related to EXPORT NODE

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

- EXPORT NODE (Export node definitions to sequential media)
You can export node definitions or file data from a server to sequential media for later importing to another server.
- EXPORT NODE (Export node definitions or file data directly to another server)
Use this command to export client node definitions or file data directly to another server for immediate import.

EXPORT NODE (Export node definitions to sequential media)

You can export node definitions or file data from a server to sequential media for later importing to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-EXPort Node----->
| .-,-----|. |
| V          | |
|---node_name-+-'|
>----->
| .-,-----|. |
| V          | |
|'-FILESpace-----file_space_name-+-'|
>----->
| .-,-----|. |
| V          | |
|'-FSID-----file_space_ID-+-'|
>----->
| .-,-----|. |
| V          | |
|'-UNIFILESpace-----file_space_name-+-'|
>----->
| .-,-----|. |
| V          | |
|'-DOfains-----domain_name-+-'|
|
|'-FILEData-----None-----|.
>----->
|'-FILEData-----+All-----+|
|          +-None-----+
|          +-ARchive-----+
|          +-Backup-----+
|          +-BACKUPActive-+
|          +-ALLActive----+
|          '-SPacemanaged-'
|
|'-Preview-----No-----|.
>----->
|          (1) (2) |
|'-Preview-----+No--+|
|          '-Yes-'
>----->
|          (1) |
|'-DEVclass-----device_class_name-'
|
|'-Scratch-----Yes-----|.
>----->
|          (2) |
|'-Scratch-----+Yes--+|
|          '-No--'
>----->
|          (2) .-,-----|. |
|          V          | |
|'-VOLumenames-----+---volume_name-+-+-'|
|          '-FILE:--file_name-'
>----->
|'-USEDVolumelist-----file_name-'
```

```

>----->
|          .-FROMTime----00:00:00-. |
'-FROMDate----date-----+'
|          '-FROMTime----time-----'

>----->
|          .-TOTime----23:59:59-. |
'-TODate----date-----+'
|          '-TOTime----time-----'

.-ENCryptionstrength----AES-----
>----->
'-ENCryptionstrength----+AES-+-'
|          '-DES-'

.-ALLOWSHREDdable----No-----
>----->
'-ALLOWSHREDdable----+No--+-'
|          '-Yes-'

```

Notes:

1. If PREVIEW=NO, a device class must be specified.
2. If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

Parameters

node_name

Specifies the client node names for which information is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. For each node entered, all file spaces in the file space, FSID, and Unicode enabled lists are searched.

Restriction: If you use wildcard characters to specify a pattern for node names, the server does not report the node names or patterns that do not match any entries in the database. Check the summary statistics in the activity log to verify that the server exported all intended nodes.

FILESpace

Specifies the file spaces for which data is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Restriction: If a file space is specified, Unicode enabled file spaces are not exported.

FSID

Specifies the file spaces by using their file space IDs (FSIDs). The server uses the FSIDs to find the file spaces to export. To find the FSID for a file space, use the QUERY FILESPACE command. Separate multiple file space IDs with commas and no intervening spaces. This parameter is optional.

UNIFILESpace

Specifies the file spaces that are known to the server as Unicode enabled. The server converts the names that you enter from the server code page to the UTF-8 code page to find the file spaces to export. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. A wildcard character can be used to specify a name. This parameter is optional.

DOmains

Specifies the policy domains from which nodes are to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. If you specify domains, a node is exported only if it belongs to one of the specified domains. You can use wildcard characters to specify a name.

FILEData

Specifies the type of files that are to be exported for all nodes that are being exported to the server. This parameter is optional. The default value is NONE.

Note: If you are exporting a node that has group data, data that is not a part of the target objects might be exported. An example of group data is virtual machine data or system state backup data. For example, if FILEDATA=BACKUPACTIVE when the FROMDATE or TODATE parameters are specified, it is possible to include inactive backup data. The incremental backup processing for the data can cause extra files that do not meet the filtering criteria to be exported.

If you are exporting to sequential media: the device class that is used by the file data is determined by the device class for the storage pool. If it is the same device class that is specified in this command, two drives are needed to export node information. The mount limit for the device class must be at least 2.

Important: If client nodes registered as TYPE=SERVER are being exported, specify ALL, ARCHIVE, or ALLACTIVE.

The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies. This parameter supports the following values:

ALL

The server exports all backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect™ for Space Management client.

None

The server does not export files, only node definitions.

ARchive

The server exports only archived files.

Backup

The server exports only backup versions, whether active or inactive.

BACKUPActive

The server exports only active backup versions. These active backup versions are the active versions in the IBM Spectrum Protect database at the time that the EXPORT command is issued.

ALLActive

The server exports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client. The active backup versions are the active versions in the IBM Spectrum Protect database at the time that the EXPORT command is issued.

SPacemanaged

The server exports only files that were migrated by an IBM Spectrum Protect for Space Management client.

Preview

Specifies whether to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data would be transferred so that you can determine how many volumes are required. This parameter supports the following values:

No

Specifies that the node information is to be exported. If you specify this value, you must also specify a device class.

Yes

Specifies that the operation is previewed but not completed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

This parameter is optional. The default value is NO.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the export runs, IBM Spectrum Protect cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER.

Scratch

Specifies whether scratch volumes can be used. The default value is YES. You can specify one of the following values:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No

Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumenames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

You can specify one of the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1-6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: <div style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;">AIX Linux /imdata/mt1.</div> <div style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;">Windows d:\program files\tivoli\tsm\data1.dsm.</div>
<div style="background-color: #f0f0f0; padding: 2px; display: inline-block;">AIX Linux Windows</div> REMOVABLEFILE	<div style="background-color: #f0f0f0; padding: 2px; display: inline-block;">AIX Linux Windows</div> 1-6 alphanumeric characters.
SERVER	1-250 alphanumeric characters.

USEDVolumelist

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. IBM Spectrum Protect ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, IBM Spectrum Protect exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- IBM Spectrum Protect ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects that are inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up after the TODATE or TOTIME parameters can be exported. An example of group data is virtual machine data or system state backup data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted 10 days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. IBM Spectrum Protect ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. An example of group data on the node is virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported so that there is a consistent image for the backup data.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW -02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. IBM Spectrum Protect ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. You can specify one of the following values:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREddable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter supports the following values:

No

Specifies that data is not exported from a storage pool that enforces shredding.

Yes

Specifies that data can be exported from a storage pool that enforces shredding. The data on the export media is not shredded.

This parameter is optional. The default value is NO.

Example: Export client node information to specific tape volumes

From the server, export client node information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be used by a device that is assigned to the MENU1 device class.

```
export node devclass=menu1 volumenames=tape01,tape02,tape03
```

Example: Export client node information by using the FSID

From the server, use the FSID to export active backup versions of file data for client node JOE to tape volume TAPE01. To determine the FSID, first issue a QUERY FILESPACE command.

1. To determine the FSID, issue a QUERY FILESPACE command.

```
query filespace joe
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity (MB)	Pct Util
JOE	\\joe\c\$	1	WinNT	NTFS	Yes	2,502.3	75.2
JOE	\\joe\d\$	2	WinNT	NTFS	Yes	6,173.4	59.6

2. Export the active backup versions of file data and specify that the tape volume is used by a device that is assigned to the MENU1 device class.

```
export node joe fsid=1,2 filedata=backupactive devclass=menu1  
volumenames=tape01
```

Example: Export client node information to tape volumes listed in a file

From the server, export client node information to tape volumes that are listed in the following file:

- **AIX** TAPEVOL
- **Linux** TAPEVOL
- **Windows** TAPEVOL.DATA

The file contains the following lines:

```
TAPE01  
TAPE02  
TAPE03
```

Specify that the tape volumes be used by a device that is assigned to the MENU1 device class. Issue the following command:

```
AIX Linux  
export node devclass=menu1 volumenames=file:tapevol
```

```
Windows  
export node devclass=menu1 volumenames=file:tapevol.data
```

EXPORT NODE (Export node definitions or file data directly to another server)

Use this command to export client node definitions or file data directly to another server for immediate import.

Important: You cannot export nodes of type NAS. Export processing excludes these nodes.

You can suspend and restart a server-to-server export operation that has a FILEDATA value other than NONE. The server saves the state and status of the export operation so that it can be restarted from the point at which the operation failed or was suspended. The export operation can be restarted later by issuing the RESTART EXPORT command.

Important: An export operation is suspended when any of the following conditions are detected:

- A SUSPEND EXPORT command is issued for the running export operation
- Segment preemption - the file that is being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered

Issue the QUERY EXPORT command to display information on any running and suspended export operations.

The export operation cannot be restarted if the export operation fails before transmitting the eligible node and file space definitions to the target server. You must reenter the command to begin a new export operation.

You can issue a QUERY PROCESS command from the target server to monitor the progress of the import operation. Issue the QUERY EXPORT command to list all restartable server-to-server export operations. See EXPORT ADMIN (Export administrator information) for a list of restrictions that apply to the export function.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
.-*-----.  
>>-EXPort Node--+----->  
      '-node_name-'  
  
>--+----->  
      '-FILESpace----file_space_name-'  
  
>--+----->  
      '-FSID----file_space_ID-'  
  
>--+----->  
      '-UNIFILESpace----file_space_name-'  
  
>--+----->  
      '-DOfains----domain_name-'  
  
. -FILEData----None-----.  
>--+----->  
      '-FILEData----+All-----+'  
                +-None-----+  
                +-ARchive-----+  
                +-Backup-----+  
                +-BACKUPActive-+  
                +-ALLActive----+  
                '-SPacemanaged-'  
  
>--+----->  
      |                .-FROMTime----00:00:00-. |  
      '-FROMDate----date--+-----+'  
                '-FROMTime----time-----'  
  
>--+----->  
      |                .-TOTime----23:59:59-. |  
      '-TODate----date--+-----+'  
                '-TOTime----time-----'  
  
>--+----->  
      '-EXPORTIDentifier----export_identifier-'
```



```

.-PREVIEWImport-----No-----.
>-----+-----+-----+----->
'-TOServer---servername-' '-PREVIEWImport-----+No--+-'
                                     '-Yes-'

.-MERGEfilespace-----No-----.
>-----+-----+-----+----->
'-MERGEfilespace-----+No--+-'
                                     '-Yes-'

.-Replacedefs-----No-----.
>-----+-----+-----+----->
'-Replacedefs-----+No--+-'
                                     '-Yes-'

.-PROXynodeassoc-----No-----.
>-----+-----+-----+----->
'-PROXynodeassoc-----+No--+-'
                                     '-Yes-'

.-ENCryptionstrength-----AES-----.
>-----+-----+-----+----->
'-ENCryptionstrength-----+AES-+-'
                                     '-DES-'

.-ALLOWSHREddable-----No-----.
>-----+-----+-----+----->>
'-ALLOWSHREddable-----+No--+-'
                                     '-Yes-'

```

Parameters

node_name

Specifies the client node names for which information is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. For each node entered, all file spaces in the file space, FSID, and Unicode enabled lists are searched.

Restriction: If you specify a list of node names or node patterns, the server does not report the node names or node patterns that do not match any of the entries in the database. Check the summary statistics in the activity log to verify that the server exported all intended nodes.

FILESpace

Specifies the file spaces for which data is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Restriction: If a file space is specified, no Unicode enabled file spaces are exported.

FSID

Specifies the file spaces by using their file space IDs (FSIDs). The server uses the FSIDs to find the file spaces to export. To find the FSID for a file space, use the QUERY FILESPACE command. Separate multiple file space IDs with commas and no intervening spaces. This parameter is optional.

UNIFILESpace

Specifies the file spaces that are known to the server to be Unicode enabled. The server converts the names that you enter from the server code page to the UTF-8 code page to find the file spaces to export. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. A wildcard character can be used to specify a name. This parameter is optional.

Domains

Specifies the policy domains from which nodes are exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. If you specify domains, IBM Spectrum Protect™ exports a node only if it belongs to one of the specified domains. You can use wildcard characters to specify a name.

FILEData

Specifies the type of files to export for all nodes. This parameter is optional. The default value is NONE.

Note: If you are exporting a node that has group data, data that is not a part of the target objects might be exported. An example of group data is virtual machine data or system state backup data. For example, if FILEDATA=BACKUPACTIVE when the FROMDATE or TODATE parameters are specified, it is possible to include inactive backup data. The incremental backup processing for the data can cause extra files that do not meet the filtering criteria to be exported.

If you are exporting to sequential media, the device class that is used by the file data is determined by the device class for the storage pool. If it is the same device class that is specified in this command, IBM Spectrum Protect requires two drives to export node information. The mount limit for the device class must be at least 2.

Important: If you export client nodes that are registered as TYPE=SERVER, specify ALL, ARCHIVE, or ALLACTIVE.
 The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies.
 The values are as follows:

ALL

The server exports all backup versions of files, all archived files, and all files that are migrated by an IBM Spectrum Protect for Space Management client.

None

The server does not export files, only node definitions.

ARChive

The server exports only archived files.

Backup

The server exports only backup versions, whether they are active or inactive.

BACKUPActive

The server exports only active backup versions. These active backup versions are the active versions in the IBM Spectrum Protect database at the time that the EXPORT command is issued.

ALLActive

The server exports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client. The active backup versions are the active versions in the IBM Spectrum Protect database at the time that the EXPORT command is issued.

SPacemanaged

The server exports only files that were migrated by an IBM Spectrum Protect for Space Management client.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. IBM Spectrum Protect ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, IBM Spectrum Protect exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files that are stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- IBM Spectrum Protect ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects that are inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up after the TODATE or TOTIME parameters can be exported. An example of group data is virtual machine data or system state backup data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted 10 days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. IBM Spectrum Protect ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. An example of group data on the node is virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported so that there is a consistent image for the backup data.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME+=02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW -02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. IBM Spectrum Protect ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

TOServer

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the DEFINE SERVER command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify TOSERVER, you cannot specify the DEVCLASS, VOLUMENAMES, and SCRATCH, USEDVOLUMELIST, and PREVIEW parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is NO.

Valid values are:

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

No

Specifies that you want the data to be imported on the target server without previewing the results.

MERGEfilespace

Specifies whether IBM Spectrum Protect merges client files into existing file spaces on the target server (if they exist), or if IBM Spectrum Protect generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No

Specifies that IBM Spectrum Protect generates a new file space name for imported data on the target server if file spaces with the same name exists.

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No

Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

PROXynodeassoc

Specifies if proxy node associations are exported. This parameter is optional. The default value is NO.

ENCCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. You can specify one of the following values:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREDdable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the server does not export data from a storage pool that enforces shredding.

Yes

Specifies that the server does export from a storage pool that enforces shredding. The data on the export media is not shredded.

Restriction: After an export operation finishes identifying files for export, any changes to the storage pool ALLOWSHREDABLE value is ignored. An export operation that is suspended retains the original ALLOWSHREDABLE value throughout the operation. You might want to consider canceling your export operation if changes to the storage pool ALLOWSHREDABLE value jeopardize the operation. You can reissue the export command after any needed cleanup.

EXPORTIdentifier

This optional parameter specifies the name that you select to identify this export operation. If you do not specify an identifier name, the server generates one for you. The export identifier name cannot be more than 64 characters, cannot

contain wildcard characters, and is not case-sensitive. You can use the identifier name to reference export operations in the QUERY EXPORT, SUSPEND EXPORT, RESTART EXPORT, or CANCEL EXPORT commands.

Restriction: You must specify the TOSERVER parameter if you are specifying the EXPORTIDENTIFIER parameter. EXPORTIDENTIFIER is ignored if FILEDATA=NONE.

Example: Export client node information and all client files

To export client node information and all client files for NODE1 directly to SERVERB, issue the following command:

```
export node node1 filedata=all toserver=serverb
```

Example: Export client node information and all client files for a specific date range

To export client node information and all client files for NODE1 directly to SERVERB between February 1, 2009 and today.

```
export node node1 filedata=all toserver=serverb
fromdate=02/01/2009 todate=today
```

Example: Export client node information and all client files for a specific date and time range

To export client node information and all client files for NODE1 directly to SERVERB from 8:00 AM on February 1, 2009 until today at 8:00 AM, issue the following command:

```
export node node1 filedata=all toserver=serverb
fromdate=02/01/2009 fromtime=08:00:00
todate=today totime=08:00:00
```

Example: Export client node information and all client files for the past three days

To export client node information and all client files for NODE1 directly to SERVERB for the past three days, issue the following command:

```
export node node1 filedata=all toserver=serverb
fromdate=today -3
```

EXPORT POLICY (Export policy information)

Use this command to export policy information from an IBM Spectrum Protect™ server to sequential media or directly to another server for immediate import. When a policy is exported by using the EXPORT POLICY command, the active data pool information in the domain is not exported.

The server exports policy information, such as:

- Policy domain definitions
- Policy set definitions, including the active policy set
- Management class definitions, including the default management class
- Backup copy group and archive copy group definitions
- Schedule definitions for each policy domain
- Client node associations, if the client node exists on the target server

You can use the QUERY ACTLOG command to view the status of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If you export policy information to sequential media and the background process is canceled, the sequential media that is holding the exported data is incomplete and must not be used to import data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, use the QUERY PROCESS command.

The following restrictions apply to the export function:

- Export operations from a later version and release to an earlier version and release is not supported.
- Export operations between servers that are at the same version and release but with different fix packs might fail. For example, you cannot export from a V7.1.3 server to a V7.1.1 or earlier server.

- Exported data from a server with retention protection enabled is not protected by retention when it is imported to another server.
- Export processing excludes nodes of type network-attached storage (NAS).
- Exporting data to a Centera device class or importing data from a Centera device class is not supported. However, files that are stored in Centera storage pools can be exported and files that must be imported can be stored on a Centera storage device.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The EXPORT POLICY command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

Table 1. Commands related to EXPORT POLICY

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT POLICY	Restores policy information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

- EXPORT POLICY (Export policy information to sequential media)
Use this command to export policy information from an IBM Spectrum Protect server to sequential media for later import to another server.
- EXPORT POLICY (Export a policy directly to another server)
Use this command to export policy information directly to another server on the network. This results in an immediate import on the target server.

EXPORT POLICY (Export policy information to sequential media)

Use this command to export policy information from an IBM Spectrum Protect™ server to sequential media for later import to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-EXPort Policy-----*----->
| .- ,----- . |
| V          | |
| ---domain_name--- |
|-----|

.-Preview-----No-----
>---+-----+----->
|          (1) (2)          |
|-----|

```


You can specify one of the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1-6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: AIX Linux /imdata/mt1. Windows d:\program files\tivoli\tsm\data1.dsm.
AIX Linux Windows REMOVABLEFILE	AIX Linux Windows 1-6 alphanumeric characters.
SERVER	1-250 alphanumeric characters.

USEDVolumelist

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

Example: Export policy information to specific tape volumes

From the server, export policy information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class.

```
export policy devclass=menu1  
volumenames=tape01,tape02,tape03
```

Example: Export policy information to tape volumes listed in a file

From the server, export policy information to tape volumes that are listed in the following file:

- **AIX** | **Linux** TAPEVOL
- **Windows** TAPEVOL.DATA

This file contains the following lines:

```
TAPE01  
TAPE02  
TAPE03
```

Specify that these tape volumes be used by a device that is assigned to the MENU1 device class. Issue the following command:

```
AIX | Linux  
export policy devclass=menu1 volumenames=file:tapevol  
  
Windows  
export policy devclass=menu1 volumenames=file:tapevol.data
```

EXPORT POLICY (Export a policy directly to another server)

Use this command to export policy information directly to another server on the network. This results in an immediate import on the target server.

To export policy information directly to SERVERB, issue the following command:

```
export policy replacedefs=yes toserver=othersrv
```

EXPORT SERVER (Export server information)

Use this command to export all or part of the server control information and client file data (if specified) from the server to sequential media.

When you export server information to sequential media, you can later use the media to import the information to another server with a compatible device type.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect™ server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the IMPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

You also have the option of processing an export operation directly to another server on the network. This results in an immediate import process without the need for compatible sequential device types between the two servers.

You can export the following types of server information by issuing the EXPORT SERVER command:

- Policy domain definitions
- Policy set definitions
- Management class and copy group definitions
- Schedules defined for each policy domain
- Administrator definitions
- Client node definitions

You can optionally export the following types of data:

- File space definitions
- Access authorization information that pertains to the file spaces exported
- Backed-up, archived, and files that were migrated by an IBM Spectrum Protect for Space Management client

This command generates a background process that can be canceled by the CANCEL PROCESS command. If you export server information to sequential media, and the background process is canceled, the sequential media holding the exported data are incomplete and should not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details.

Issue the QUERY PROCESS command from the target server to monitor the progress of the import operation. Issue the QUERY EXPORT command to list all server-to-server export operations (that have a FILEDATA value other than NONE) that are running or suspended.

You can use the QUERY ACTLOG command to view the actual status information which indicates the size and the success or failure of the export operation.

The following restrictions apply to the export function:

- Export operations from a later version and release to an earlier version and release is not supported.
- Export operations between servers that are at the same version and release but with different fix packs might fail. For example, you cannot export from a V7.1.3 server to a V7.1.1 or earlier server.
- Exported data from a server with retention protection enabled is not protected by retention when it is imported to another server.
- Export processing excludes nodes of type network-attached storage (NAS).
- Exporting data to a Centera device class or importing data from a Centera device class is not supported. However, files that are stored in Centera storage pools can be exported and files that must be imported can be stored on a Centera storage device.
- The EXPORT NODE and EXPORT SERVER commands do not export data from a shred pool unless you explicitly allow it by setting the ALLOWSHREDDABLE parameter to the YES value. If this value is specified, and the exported data includes data from shred pools, that data cannot be shredded. A warning is not issued if the export operation includes data from shred pools.

- Incrementally exporting or importing the following types of client data to another IBM Spectrum Protect server is not supported:
 - VMware backups where full plus incremental backups need to be periodically, incrementally transferred to another server
 - Backups groups where full plus differential backups must be periodically, incrementally transferred to another server
 - Windows System State data that is periodically, incrementally transferred to another server

Full export or import of this data to a new file system on the target is supported by exporting the entire file space that contains the data. The export must not use the FILEDATA=ALLACTIVE, FROMDATE, TODATE, or MERGEFILESPPACES parameters.

Using node replication to incrementally transfer this type of client data between two servers is optimal.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The EXPORT SERVER command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

Table 1. Commands related to EXPORT SERVER

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVE DATA	Copies active backup data.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

- EXPORT SERVER (Export a server to sequential media)
You can export all or part of the server control information and client file data from a server to sequential media so that this information can be imported to another server.
- EXPORT SERVER (Export server control information and client file data to another server)
Use this command to export all or part of the server control information and client file data directly to another server on the network. This results in an immediate import on the target server.

EXPORT SERVER (Export a server to sequential media)

You can export all or part of the server control information and client file data from a server to sequential media so that this information can be imported to another server.

Specifies the type of files that are exported for all nodes that are defined to the server. This parameter is optional. The default value is NONE.

If you are exporting to sequential media, the device class to access the file data is determined by the device class for the storage pool. If it is the same device class that is specified in this command, two drives are needed to export server information. The mount limit for the device class must be set to at least 2.

The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies. The following values are available:

ALL

IBM Spectrum Protect™ exports all backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client.

None

IBM Spectrum Protect does not export files, only definitions.

ARchive

IBM Spectrum Protect exports only archived files.

Backup

IBM Spectrum Protect exports only backup versions, whether the versions are active or inactive.

BACKUPActive

IBM Spectrum Protect exports only active backup versions.

ALLActive

IBM Spectrum Protect exports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client.

SPacemanaged

IBM Spectrum Protect exports only files that were migrated by an IBM Spectrum Protect for Space Management client.

Preview

Specifies whether you want to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data are transferred so that you can determine how many volumes are required. This parameter supports the following values:

No

Specifies that the server information is to be exported. If you specify this value, you must also specify a device class.

Yes

Specifies that the operation is previewed but not completed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

This parameter is optional. The default value is NO.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the export runs, IBM Spectrum Protect cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER.

Scratch

Specifies whether scratch volumes can be used. The default value is YES. You can specify one of the following values:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No

Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumenames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

You can specify one of the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify						
Tape	1-6 alphanumeric characters.						
FILE	Any fully qualified file name string. For example: <table border="1"><tr><td>AIX</td><td>Linux</td></tr></table> /imdata/mt1. <table border="1"><tr><td>Windows</td></tr></table> d:\program files\tivoli\tsm\data1.dsm.	AIX	Linux	Windows			
AIX	Linux						
Windows							
<table border="1"><tr><td>AIX</td><td>Linux</td><td>Windows</td></tr></table> REMOVABLEFILE	AIX	Linux	Windows	<table border="1"><tr><td>AIX</td><td>Linux</td><td>Windows</td></tr></table> 1-6 alphanumeric characters.	AIX	Linux	Windows
AIX	Linux	Windows					
AIX	Linux	Windows					
SERVER	1-250 alphanumeric characters.						

USEDVolumelist

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. IBM Spectrum Protect ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, IBM Spectrum Protect exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- IBM Spectrum Protect ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. IBM Spectrum Protect ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. An example of group data on the node is virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup

processing can cause extra files that do not meet the filtering criteria to be exported so that there is a consistent image for the backup data.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or - HH:MM	The current time minus hours and minutes specified	NOW -02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. IBM Spectrum Protect ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value. Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. You can specify one of the following values:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREDdable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that data is not exported from a storage pool that enforces shredding.

Yes

Specifies that data can be exported from a storage pool that enforces shredding. The data on the export media is not shredded.

Example: Export a server to specific tape volumes

From the server, export server information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class.

```
export server devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Export a server to tape volumes listed in a file

From the server, export server information to tape volumes that are listed in the following file:

- **AIX** | **Linux** TAPEVOL
- **Windows** TAPEVOL.DATA

The file contains the following lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that the tape volumes be used by a device that is assigned to the MENU1 device class. Issue the following command:

```
AIX | Linux
export server devclass=menu1 volumenames=file:tapevol

Windows
export server devclass=menu1 volumenames=file:tapevol.data
```

EXPORT SERVER (Export server control information and client file data to another server)

Use this command to export all or part of the server control information and client file data directly to another server on the network. This results in an immediate import on the target server.

Server-to-server export operations that have a FILEDATA value other than NONE can be restarted after the operation is suspended. The server saves the state and status of the export operation so that it may be restarted from the point at which the operation failed or was suspended. The export operation can be restarted at a later date by issuing the RESTART EXPORT command. These export operations can be manually suspended as well as restarted. Therefore, if an export fails, it is automatically suspended if it has completed the transmitting definitions phase.

An export operation is suspended when any of the following conditions is detected:

- A SUSPEND EXPORT command is issued for the running export operation
- Segment preemption - the file being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered

The export operation cannot be restarted if the export operation fails prior to transmitting the eligible node and filespace definitions to the target server. You must reenter the command to begin a new export operation.

Issue the QUERY PROCESS command from the target server to monitor the progress of the import operation. Issue the QUERY EXPORT command to list all server-to-server export operations (that have a FILEDATA value other than NONE) that are running or suspended. See EXPORT ADMIN (Export administrator information) for a list of restrictions that apply to the export function.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-FILEData----None-----.
>>-EXPort Server-----+----->
      '-FILEData----+All-----+'
                        +-None-----+
                        +-ARchive-----+
                        +-Backup-----+
                        +-BACKUPActive--+
                        +-ALActive-----+
                        '-SPacemanaged-'

>-----+----->
|                                     .-FROMTime----00:00:00-. |
'-FROMDate---date-----+-----+'
      '-FROMTime----time-----'

>-----+----->
|                                     .-TOTime----23:59:59-. |
'-TODate---date-----+-----+'
      '-TOTime----time-----'

>-----+----->
'-EXPORTIDentifier----export_identifier-'

                                     .-PREVIEWImport----No-----.
>-----+-----+-----+----->
'-TOServer---servername-' '-PREVIEWImport----+No--+-'
                                     '-Yes-'

      .-MERGEfilespace----No-----.
>-----+-----+----->
'-MERGEfilespace----+No--+-'
                                     '-Yes-'

      .-Replacedefs----No-----.
>-----+-----+----->
'-Replacedefs----+No--+-'
                                     '-Yes-'

      .-PROXynodeassoc----No-----.
>-----+-----+----->
'-PROXynodeassoc----+No--+-'
                                     '-Yes-'

      .-ENCryptionstrength----AES-----.
>-----+-----+----->
'-ENCryptionstrength----+AES--+-'
                                     '-DES-'

      .-ALLOWSHREddable----No-----.
>-----+-----+-----><
'-ALLOWSHREddable----+No--+-'
                                     '-Yes-'

```

Parameters

FILEData

Specifies the type of files to export for all nodes defined to the server. This parameter is optional. The default value is NONE.

If you are exporting to sequential media: The device class to access the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, IBM Spectrum Protect™ requires two drives to export server information. You must set the mount limit for the device class to at least 2.

The following descriptions mention active and inactive backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies.

The values are:

ALL

IBM Spectrum Protect exports all backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client.

None

IBM Spectrum Protect does not export files, only definitions.

ARchive

IBM Spectrum Protect exports only archived files.

Backup

IBM Spectrum Protect exports only backup versions, whether they are active or inactive.

BACKUPActive

IBM Spectrum Protect exports only active backup versions.

ALLActive

IBM Spectrum Protect exports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client.

SPacemanaged

IBM Spectrum Protect exports only files that were migrated by an IBM Spectrum Protect for Space Management client.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. IBM Spectrum Protect ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported, so that there is a consistent image for the backup data.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, IBM Spectrum Protect exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for

selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- IBM Spectrum Protect ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. IBM Spectrum Protect ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If you have group data on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can also be exported. An example of group data on the node is virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported so that there is a consistent image for the backup data.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
-------	-------------	---------

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW -02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. IBM Spectrum Protect ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value. Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

TOServer

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the DEFINE SERVER command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify TOSERVER, you cannot specify the DEVCLASS, VOLUMENAMES, and SCRATCH, USEDVOLUMELIST, and PREVIEW parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is NO.

Valid values are:

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

No

Specifies that you want the data to be imported on the target server without previewing the results.

MERGEfilespace

Specifies whether IBM Spectrum Protect merges client files into existing file spaces on the target server (if they exist), or if IBM Spectrum Protect generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No

Specifies that IBM Spectrum Protect generates a new file space name for imported data on the target server if file spaces with the same name exists.

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No

Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

PROXynodeassoc

Specifies if proxy node associations are exported. This parameter is optional. The default value is NO.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. You can specify one of the following values:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREDdable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the server does not allow data to be exported from a storage pool that enforces shredding.

Yes

Specifies that the server allows data to be exported from a storage pool that enforces shredding. The data on the export media will not be shredded.

Important: After an export operation finishes identifying files for export, any changes to the storage pool ALLOWSHREDABLE value is ignored. An export operation that is suspended retains the original ALLOWSHREDABLE value throughout the operation. You might want to consider canceling your export operation if changes to the storage pool ALLOWSHREDABLE value jeopardize the operation. You can reissue the export command after any needed cleanup.

EXPORTIDentifier

This optional parameter specifies the name that you selected to identify this export operation. If you do not specify a command name, the server generates one for you. The export identifier name cannot be more than 64 characters, cannot contain wildcard characters, and is not case sensitive. You can use the identifier name to reference export operations in the QUERY EXPORT, SUSPEND EXPORT, RESTART EXPORT, or CANCEL EXPORT commands. EXPORTIDENTIFIER is ignored if FILEDATA=NONE or if PREVIEWIMPORT=YES.

If you are specifying the EXPORTIDENTIFIER parameter, you must specify the TOSERVER parameter.

Example: Export server information directly to another server

To export server information directly to SERVERB, issue the following command.

```
export server filedata=all toserver=serverb
```

Example: Export server information directly to another server using a date range

To export directly to SERVERB between February 1, 2009 and today, issue the following command.

```
export server filedata=all toserver=serverb
fromdate=02/01/2009 todate=today
```

Example: Export server information and client file data directly to another server using a date and time range

To export directly to SERVERB from 8:00 a.m. on February 1, 2009 until today at 8:00 a.m., issue the following command.

```
export server filedata=all toserver=serverb
fromdate=02/01/2009 fromtime=08:00:00
todate=today totime=08:00:00
```

EXTEND DBSPACE (Increase space for the database)

Use this command to increase space for the database by adding directories for the database to use.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

When you issue the EXTEND DBSPACE command, directories are added to the database. With the default parameter settings, data is redistributed across all database directories, and storage space is reclaimed. This action improves parallel I/O performance and makes the new directory space available for immediate use.

If you do not want to redistribute data when you add new directories, you can specify `RECLAIMSTORAGE=NO`. If you specify `NO` for this parameter, all space in existing directories is filled before new directories are used. You can redistribute data and reclaim space later, but you must complete the manual procedure for this task by using DB2 commands.

Restriction: Redistribution of data and reclaiming of space as part of an operation to extend database space works only with DB2 Version 9.7 or later table spaces. The table spaces are created when you format a new IBM Spectrum Protect™ Version 6.2 or later server. If you upgraded or restored your IBM Spectrum Protect server from V6.1, you cannot redistribute data or reclaim space. You must issue the EXTEND DBSPACE command with `RECLAIMSTORAGE=NO`.

Important: The redistribution process uses considerable system resources, so ensure that you plan ahead when you want to add space to the database. Review the following guidelines:

- Complete the process when the server is not handling a heavy workload.
- The time that is required to redistribute data and reclaim space might vary. It is affected by factors such as the file system layout, the ratio of new paths to existing storage paths, server hardware, and concurrent operations. To get a rough estimate, you can try the operation with a small IBM Spectrum Protect database on a lab system. Use your results as a reference to estimate the time that is required for the procedure.
- Do not interrupt the redistribution process. If you try to stop it, for example, by halting the process that is completing the work, you must stop and restart the DB2® server. When the server is restarted, it will go into crash recovery mode, which takes several minutes, after which the redistribution process resumes.

After an operation to extend the database space is complete, halt and restart the server to fully use the new directories. If the existing database directories are nearly full when a new directory is added, the server might encounter an out of space condition (reported in the `db2diag.log`). You can fix the out of space condition by halting and restarting the server.

Syntax

```
      .-,------.
      v          |
>>-EXTend DBSpace---db_directory+----->

      .-REclaimstorage---Yes----- .-Wait-----No-----
>-+-----+-----+-----+-----><
      '-REclaimstorage---+No--+-' '-Wait---+No--+-'
          '-Yes-'          '-Yes-'
```


Parameters

db_directory (Required)

Specifies the directories for database storage. The directories must be empty and accessible by the user ID of the database manager. A directory name must be a fully qualified name and cannot exceed 175 characters in length. Enclose the name in quotation marks if it contains embedded blanks, an equal sign, or other special characters. If you are specifying a list of directories for database storage, the maximum length of the list can be 1400 characters.

Windows Restriction: You cannot specify Universal Naming Convention (UNC) paths.

Tip: Specify directories that are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

REClaimstorage

Specifies whether data is redistributed across newly created database directories and space is reclaimed from the old storage paths. This parameter is optional. The default value is Yes.

Unless you specify `WAIT=YES`, the operation is completed as a background process.

Yes

Specifies that data is redistributed so that new directories are available for immediate use.

Important: The redistribution process uses considerable system resources so ensure that you plan ahead.

After the process starts, messages are issued to inform you about the progress. You can use the `QUERY PROCESS` command to monitor the operation. To cancel the process, you can use the `CANCEL PROCESS` command, but if a data redistribution operation is in progress, it completes before the process is stopped.

No

Specifies that data is not redistributed across database directories and storage space is not reclaimed when space is added for the database.

Wait

Specifies whether this command is processed in the background or foreground.

No

Specifies background processing. The default is NO.

Yes

Specifies foreground processing.

AIX | **Linux** You cannot specify YES from the server console.

AIX | **Linux**

Example: Add directories to the storage space for the database, redistribute data, and reclaim storage

Add two directories (`/tsm_db/stg1` and `tsm_db/stg2`) under the `/tsm_db` directory to the storage space for the database. Issue the command:

```
extend dbspace /tsm_db/stg1,/tsm_db/stg2
```

Windows

Example: Add drives to the storage space for the database, redistribute data, and reclaim storage

Add drives D and E to the storage space for the database. Issue the command:

```
extend dbspace D:,E:
```

Related commands

Table 1. Commands related to EXTEND DBSPACE

Command	Description
DSMSERV EXTEND DBSPACE	Adds directories to increase space for use by the database.

Command	Description
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

Related tasks:

Managing inventory capacity

GENERATE commands

Use the GENERATE commands for backup sets for a selected filesystem or client node.

- GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)
- GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)
- AIX Linux Windows GENERATE DEDUPSTATS (Generate data deduplication statistics)

GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)

Use this command to generate a backup set for a Backup-Archive Client node. A *backup set* is a collection of a Backup-Archive Client's active backed up data, which is stored and managed as a single object, on specific media, in server storage. Although you can create a backup set for any client node, a backup set can be used only by a Backup-Archive Client.

Restriction: A backup set in "deduplication format" has that designation as a result of a GENERATE BACKUPSET command with at least one of the following specifications:

- Includes a node at Backup-Archive Client Version 6.1.x (at least V6.1.0 but less than V6.2.0).
- Includes a node that has one or more nodes that are authorized to act as a proxy. At least one of those proxy nodes is at Backup-Archive Client V6.1.x.

Backup sets in the deduplication format can be restored only by the V6.1.2 or later Backup-Archive Client. Backup-Archive Clients before V6.1.2 cannot restore from a backup set that is in the deduplication format.

A backup set in the "distributed deduplication format" has that designation as a result of a GENERATE BACKUPSET command with at least one of the following specifications:

- Includes a node at Backup-Archive Client level V6.2.0 or later.
- Includes a node that has one or more nodes that are authorized to act as a proxy. At least one of those proxy nodes is at Backup-Archive Client V6.2.0.

Backup sets in the distributed deduplication format can be restored only by the V6.2.0 or later Backup-Archive Client.

Restriction: You cannot generate a backup set with files that were backed up to IBM Spectrum Protect™ using NDMP. However, you can create a backup set with files that were backed up using NetApp SnapShot Difference.

The server creates copies of active versions of a client's backed up objects that are within the one-or-more file spaces specified with this command. The server then consolidates them onto sequential media. Currently, the backup object types that are supported for backup sets include directories and files only.

The backup-archive client node can restore its backup set from the server and from the media to which the backup set was written.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If the background process created by this command is canceled, the media might not contain a complete backup set. You can use the QUERY PROCESS command to show information about the background process that is created by this command.

Tip: When IBM Spectrum Protect generates a backup set, you can improve performance if the primary storage pools containing the client data are collocated. If a primary storage pool is collocated, client node data is likely to be on fewer tape volumes than it would be if the storage pool were not collocated. With collocation, less time is spent searching database entries, and fewer mount operations are required.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```

      .-,------.
      v           |
>>-GENERATE BACKUPSET-----+node_name-----+----->
      '-node_group_name-'

      .-*------.
>>-backup_set_name_prefix--+-----+----->
      | .-,------. |
      | v           | |
      |'---file_space_name+-'

      .-SCRatch---Yes-----.
>>-DEVclass---device_class_name-----+----->
      '-SCRatch---+Yes+-'
      '-No--'

>+-----+----->
|           .-,------. |
|           v           | |
|'-VOLumes-----volume_names+-'

      .-RETention---365-----.
>+-----+----->
|'-RETention---+days+-'
|           '-NOLimit-'

      .-Wait---No-----.
>+-----+-----+-----+----->
|'-DESCRIPTION---description-' |'-Wait---+No+-'
|                               |'-Yes-'

      .-NAMEType---SERVER-----.
>+-----+-----+----->
|'-NAMEType---+SERVER+-'
|           +-UNICODE+
|           '-FSID----'

      .-CODEType---BOTH-----.
>+-----+-----+----->
|'-CODEType---+UNICODE+-'
|           +-NONUNICODE+
|           '-BOTH-----'

      .-PITDate---current_date-.   .-PITTime---current_time-.
>+-----+-----+-----+----->
|'-PITDate---date-----' |'-PITTime---time-----'

      .-DATAType---FILE-----.   .-TOC---Preferred-----.
>+-----+-----+-----+----->
|           .-,------. | |'-TOC---+No-----+-'
|           v           | |           +-Preferred+
|'-DATAType---+FILE--+-' |           '-Yes-----'
|           +-IMAGE+
|           '-ALL----'

>+-----+-----+----->
|'-TOCMgmtclass---class_name-'

      .-ALLOWSHREddable---No-----.
>+-----+-----+----->>
|'-ALLOWSHREddable---+No+-'
|           '-Yes-'

```

Parameters

node_name or node_group_name (Required)

Specifies the name of the client node and node groups whose data is contained in the backup set. To specify multiple node names and node group names, separate the names with commas and no intervening spaces. You can use wildcard characters with node names but not with node group names. When multiple node names are specified, the server generates a backup set for each node and places all of the backup sets together on a single set of output volumes.

backup_set_name_prefix (Required)

Specifies the name of the backup set for the client node. The maximum length of the name is 30 characters.

When you select a name, IBM Spectrum Protect adds a suffix to construct your backup set name. For example, if you name your backup set *mybackupset*, IBM Spectrum Protect adds a unique number such as 3099 to the name. The backup set name is then identified to IBM Spectrum Protect as *mybackupset.3099*. To later show information about this backup set, you can include a wildcard with the name, such as *mybackupset.** or specify the fully qualified name, such as *mybackupset.3099*.

When multiple node names or node group names are specified, the server generates a backup set for each node or node group and places all the backup sets on a single set of output volumes. Each backup set is given the same fully qualified name consisting of the *backup_set_name_prefix* and a suffix determined by the server.

file_space_name

Specifies the names of one or more file spaces that contain the data to be included in the backup set. This parameter is optional. The file space name that you specify can contain wildcard characters. You can specify more than one file space by separating the names with commas and no intervening spaces. If you do not specify a file space, data from all the client nodes backed-up and active file spaces is included in the backup set.

For a server that has clients with support for Unicode-enabled file spaces, you can enter either a file space name or a file space ID (FSID). If you enter a file space name, you might need to have the server convert the file space name that you enter. For example, you might need to have the server convert the name that you enter from the server's code page to Unicode. See the NAMETYPE parameter for details. If you do not specify a file space name, or specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

DEVclass (Required)

Specifies the name of the device class for the volumes to which the backup set is written. The maximum length of the name is 30 characters.

Restriction: You cannot specify a device class with a device type of NAS or CENTERA.

SCRatch

Specifies whether to use scratch volumes for the backup set. If you include a list of volumes using the VOLUMES parameter, the server uses scratch volumes only if the data cannot be contained in the volumes you specify. The default is SCRATCH=YES. The values are:

YES

Specifies to use scratch volumes for the backup set.

NO

Specifies not to use scratch volumes for the backup set.

VOLumes

Specifies the names of one or more volumes that will contain the backup set. This parameter is optional. You can specify more than one volume by separating each volume with a comma, with no intervening spaces.

If you do not specify this parameter, scratch volumes are used for the backup set.

RETention

Specifies the number of days to retain the backup set on the server. You can specify an integer from 0 to 30000. The default is 365 days. The values are:

days

Specifies the number of days to retain the backup set on the server.

NOLimit

Specifies that the backup set should be retained on the server indefinitely.

If you specify NOLIMIT, the server retains the volumes containing the backup set forever, unless a user or administrator deletes the volumes from server storage.

DESCription

Specifies the description to associate with the backup set. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. The values are:

Yes

Specifies the command processes in the foreground. Messages that are created are not displayed until the command completes processing. You cannot specify WAIT=YES from the server console.

No

Specifies that the command processes in the background. Use the QUERY PROCESS command to monitor the background processing of this command.

NAMETYPE

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode-enabled file spaces. You can use this parameter for IBM Spectrum Protect clients using Windows, NetWare, or Macintosh OS X operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

Important: Use care when specifying this parameter if multiple node names are also specified. Different nodes might use the same file space ID for different file spaces, or different file space IDs for the same file space name.

Therefore, specifying a file space ID as the file space names can result in the wrong data being written to the backup set for some nodes.

CODETYPE

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name or when you do not specify any file space names. Possible values are:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

PITDATE

Specifies that files that were active on the specified date and that are still stored on the IBM Spectrum Protect server are to be included in the backup set, even if they are inactive at the time you issue the command. This parameter is optional. The default is the date on which the GENERATE BACKUPSET command is run. You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-daysor-days	The current date minus days specified	TODAY-7 or -7. To include files that were active a week ago, specify PITDATE=TODAY-7 or PITDATE=-7
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

Value	Description	Example
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

PITTime

Specifies that files that were active on the specified time and that are still stored on the IBM Spectrum Protect server are to be included in the backup set, even if they are inactive at the time you issue the command. This parameter is optional. If a PITDate was specified, the default is midnight (00:00:00); otherwise the default is the time at which the GENERATE BACKUPSET command is started. You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified PIT date	12:33:28
NOW	The current date on the specified PIT date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified PIT date	NOW+03:00 or +03:00 If you issue this command at 9:00 with PITTIME=NOW+03:00 or PITTIME=+03:00. IBM Spectrum Protect includes files that were active at 12:00 on the PIT date.

DATATYPE

Specifies that backup sets containing the specified types of data that are to be generated. This parameter is optional. The default is that file level backup sets are to be generated. To specify multiple data types, separate data types with commas and no intervening spaces.

The server generates a backup set for each data type and places all the backup sets on a single set of output volumes. Each backup set is given the same fully qualified name consisting of the *backup_set_name_prefix* and a suffix determined by the server. However, each backup set has a different data type, as shown by the QUERY BACKUPSET command. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) that have been backed up on the server are to be generated.

FILE

Specifies that a file level backup set is to be generated. File level backup sets contain files and directories that are backed up by the backup client. If no files or directories have been backed up by the backup client, a file level backup set is not generated. This is the default.

IMAGE

Specifies that an image backup set is to be generated. Image backup sets contain images that are created by the backup client BACKUP IMAGE command. Image backup sets are generated only if an image has been backed up by the backup client.

TOC

Specifies whether a table of contents (TOC) is saved for each file level backup set. Tables of contents are always saved for backup sets containing image or application data. The TOC parameter is ignored when generating image and application backup sets. A table of contents will always be generated for image and application backup sets.

Consider the following in determining whether you want to save a table of contents:

- If a table of contents is saved for a backup set, you can use the IBM Spectrum Protect web backup-archive client to examine the entire file system tree and choose files and directories to restore. To create a table of contents, you must define the TOCDESTINATION attribute in the backup copy group for the management class that is specified by

the TOCMGMTCLASS parameter. Creating a table of contents requires additional processing, storage pool space, and possibly a mount point during the backup set operation.

- If a table of contents is not saved for a backup set, you can still restore individual files or directory trees using the backup-archive client RESTORE BACKUPSET command, if you know the fully qualified name of each file or directory to be restored.

To display the contents of backup sets, you can also use the QUERY BACKUPSETCONTENTS command.

This parameter is optional. Possible values are:

No

Specifies that table of contents information is not saved for file level backup sets.

Preferred

Specifies that table of contents information should be saved for file level backup sets. This is the default. However, a backup set does not fail just because an error occurs during creation of the table of contents.

Yes

Specifies that table of contents information must be saved for each file level backup set. A backup set fails if an error occurs during creation of the table of contents.

TOCMgmtclass

Specifies the name of the management class to which the table of contents should be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. In this case, creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the specified management class.

ALLOWSHREddable

Specifies whether data from a storage pool that enforces shredding is included in the backup set. This parameter is optional. Possible values are:

No

Specifies that data from a storage pool that enforces shredding is not included in the backup set. This is the default.

Yes

Specifies that data from a storage pool that enforces shredding can be included in the backup set. The data on the backup set media will not be shredded.

Example: Generate a backup set for a file space

Generate a backup set of a file space that is called /srvr that belongs to client node JANE. Name the backup set PERS_DATA and retain it for 75 days. Specify that volumes VOL1 and VOL2 contain the data for the backup set. The volumes are to be read by a device that is assigned to the AGADM device class. Include a description.

```
generate backupset jane pers_data /srvr devclass=agadm
retention=75 volumes=vol1,vol2
description="area 51 base image"
```

Example: Generate a backup set of a Unicode-enabled file space

Generate a backup set of the Unicode-enabled file space, \\joe\c\$, that belongs to client node JOE. Name the backup set JOES_DATA. Specify that volume VOL1 contain the data for the backup set. The volume is to be read by a device that is assigned to the AGADM device class. Have the server convert the \\joe\c\$ file space name from the server code page to the UTF-8 code page.

```
generate backupset joe joes_data \\joe\c$ devclass=agadm
volumes=vol1 nametype=unicode
```

Related commands

Table 1. Commands related to GENERATE BACKUPSET

Command	Description
CANCEL PROCESS	Cancel a background server process.
COPY ACTIVATEDATA	Copies active backup data.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.

Command	Description
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
QUERY BACKUPSET	Displays backup sets.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY NODEGROUP	Displays information about node groups.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE NODEGROUP	Updates the description of a node group.

GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)

Use this command to generate a table of contents for a backup set that does not already have one. The backup-archive client uses the table of contents to display the backup set, which allows users to select individual files to be restored from the backup set.

Creating a table of contents for a backup set requires storage pool space and possibly one or more mount points during the creation operation.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```
>>-GENerate BACKUPSETTOC--node_name--backup_set_name----->
      .-DATAType-----ALL-----
>--+-----+-----+-----+-----+-----+-----+-----+-----+----->
      |           .-,-----.|
      |           V           ||
      |'-DATAType-----+FILE--+-'
      |           '-IMAGE-'
>--+-----+-----+-----+-----+-----+-----+-----+-----+-----><
      |'-TOCMgmtclass-----class_name-'
```

Parameters

node_name (Required)

Specifies the name of the client node whose data is contained in the backup set. You cannot use wildcard characters to specify a name, nor can you specify a list of client node names.

backup_set_name (Required)

Specifies the name of the backup set for the client node. You cannot use wildcard characters to specify a name, nor can you specify a list of backup set names.

DATAType

Specifies the type of data to be included in the table of contents. This parameter is optional. By default, all data is included. To specify multiple data types, separate the data types with commas and no intervening spaces. Possible values are:

ALL

Specifies that the table of contents includes all types of data (file-level, image, and application) stored in the backup set. This is the default.

FILE

Specifies that the table of contents includes only file-level data. File-level data consists of files and directories backed up by the backup-archive client. If the backup set contains no files or directories, the table of contents is not generated.

IMAGE

Specifies that the table of contents will include only image backups. Image backups consist of file system images created by the backup client BACKUP IMAGE command. If the backup set contains no image backups, the table of contents will not be generated.

TOCMgmtclass

Specifies the name of the management class to which the table of contents should be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. If you create a table of contents you must define the TOCDESTINATION attribute in the backup copy group for the specified management class.

Example: Generate a table of contents

Generate a table of contents for a backup set named PROJX_DATA that contains the data for client node GARY. The table of contents is to be bound to the default management class.

```
generate backupsettoc gary projx_data
```

Related commands

Table 1. Commands related to GENERATE BACKUPSETTOC

Command	Description
COPY ACTIVE DATA	Copies active backup data.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE NODEGROUP	Updates the description of a node group.

AIX

Linux

Windows

GENERATE DEDUPSTATS (Generate data deduplication statistics)

Use this command to generate data deduplication statistics for a directory-container storage pool or a cloud-container storage pool to determine data deduplication performance.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool.

Syntax

```
>>-GENerate DEDUPStats--pool_name----->
. ,-----
V | .-*-----
>-----+node_name-----+-----+----->
  '-node_group_name-' | . ,----- |
                        | V | |
                        +---+filespace_name+---+
                        | . ,----- |
                        | V | |
                        '-----FSID-----'

.-CODEType---==--BOTH----- . .-MAXProcess---==--4-----
>-----+-----+-----+----->
  '-CODEType---==--+-UNICODE----+' '-MAXProcess---==--number-'
    +-NONUNICODE+
    '-BOTH-----'

.-NAMEType---==--SERVER----- . .-Wait---==--No-----
>-----+-----+-----+----->
  '-NAMEType---==--+-SERVER--+-' '-Wait---==--+-No--+-'
    +-UNICODE+
    '-FSID----'

>-----+-----+-----+-----><
  '-DESCRiption---==--description-'
```

Parameters

pool_name (Required)

Specifies the name of the storage pool that is reported in the data deduplication statistics. You can specify up to 30 characters for the storage pool name. If you specify more than 30 characters, the command fails.

Restriction: You can specify only directory-container storage pools or cloud storage pools.

node_name or node_group_name (Required)

Specifies the name of the client node or defined group of client nodes that is reported in the data deduplication statistics.

You can also specify a combination of client node names and client-node group names. To specify multiple client node names or client-node group names, separate the names with commas with no intervening spaces. You can use wildcard characters with client node names but not with client-node group names. The specified value can have a maximum of 1024 characters.

filespace_name or FSID

Specifies the names of one or more file spaces for which data deduplication statistics are collected. This parameter is optional. You can use wildcard characters to specify this name. The specified value can have a maximum of 1024 characters. An asterisk is the default. You can specify one of the following values:

*

Specify an asterisk (*) to show information for all file spaces or IDs.

filespace_name

Specifies the name of the file space. You can specify more than one file space by separating the names with commas and no intervening spaces.

FSID

Specifies the name of a file space identifier. This parameter is valid for clients with file spaces that are in Unicode format. Specify more than one file space by separating the names with commas and no intervening spaces.

For clients with file spaces that are in Unicode format, you can enter either a file space name or an FSID. If you enter a file space name, the server might have to convert the file space name that you enter. For example, the server might have to convert the name that you enter from the server's code page to Unicode.

Restrictions: The following restrictions apply to file space names and FSIDs:

- You must specify a node name if you specify a file space name.
- Do not specify both file space names and FSIDs on the same command.

CODEType

Specifies what type of file spaces to include in the record. The default value is BOTH, which specifies that file spaces are included regardless of code page type. Use this parameter only when you enter an asterisk to display information about all file spaces. This parameter is optional. Specify one of the following values:

UNICODE

Include file spaces that are in Unicode format.

NONUNICODE

Include file spaces that are not in Unicode format.

BOTH

Include file spaces regardless of code page type. This is the default.

MAXPROcess

Specifies the maximum number of parallel processes to generate statistics for a container in a directory-container or cloud-container storage pool. This parameter is optional. Enter a value in the range 1 - 99. The default value is 4.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. Use this parameter when IBM Spectrum Protect™ clients have file spaces that are in Unicode format, and are on Windows, NetWare, or Macintosh OS X operating systems. This parameter is optional.

This parameter is required if you specify a node name and a file space name or FSID.

Restriction: When you specify this parameter, the file space name cannot contain an asterisk.

Specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names. This is the default.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server's code page.

Tip: Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their FSIDs.

Wait

Specifies whether the data deduplication statistics are generated in the foreground or background. This parameter is optional. You can specify one of the following values:

No

Specifies that the operation is completed in the background. You can continue with other tasks while the command is processing. Messages that are related to the background process are displayed in the activity log file or the server console, depending on where the messages are logged. This is the default value.

Yes

Specifies that the operation is completed in the foreground. It might take a long time to complete the operation. The operation must end before you can continue with other tasks. Messages are displayed in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify the WAIT=YES parameter from the server console.

DESCription

Specifies a description of the generated statistics. This parameter is optional.

Example: Generate data deduplication statistics for a file space

Generate data deduplication statistics for a file space that is called /srvr that belongs to a directory-container storage pool, POOL1, that is stored on client node NODE1.

```
generate dedupstats pool1 node1 /srvr
```

Example: Generate data deduplication statistics for a Unicode-enabled file space

Generate data deduplication statistics for a Unicode-enabled file space that is called \\abc\c\$ that belongs to client node NODE2. Convert the \\abc\c\$ file space name from the server code page to the UTF-8 code page.

```
generate dedupstats node2 \\abc\c$ nametype=unicode
```

Related commands

Table 1. Commands related to GENERATE DEDUPSTATS

Command			Description	
AIX	Linux	Windows	DELETE DEDUPSTATS	Deletes data deduplication statistics.
AIX	Linux	Windows	QUERY DEDUPSTATS	Displays data deduplication statistics.

GRANT commands

Use the GRANT command to grant appropriate privileges or access.

- GRANT AUTHORITY (Add administrator authority)
- GRANT PROXYNODE (Grant proxy authority to a client node)

GRANT AUTHORITY (Add administrator authority)

Use this command to grant an administrator one or more administrative privilege classes, and authority to access client nodes.

You cannot grant restricted privilege to an unrestricted policy or unrestricted storage administrator. You must use the REVOKE AUTHORITY command to remove the administrator's unrestricted privilege, then use this command to grant restricted privilege to the administrator.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-GRant AUTHority--admin_name----->
      .-,------.
      (1)  V          |
>>-Classes-----+--SYSTEM-----+----->
      +--Policy-----+
      +--Storage-----+
      +--Operator-----+
      '-Node--| A |-'
```

```
>----->
|      .-,------.
|      V          | |
| '-DObains-----domain_name-+-'
```

```
>----->>
|      .-,------.
|      (1)  V          | |
| '-STGpools-----pool_name-+-'
```

A

```
.-AUTHority-----Access-----
|-----+-----+-----+-----+-----|
| '-AUTHority-----+Access-+-'   '-NObde-----node_name-----'
|                               '-Owner--'
```

Notes:

1. You must specify one or more of these parameters.

Parameters

admin_name (Required)

Specifies the name of the administrator being granted an administrative privilege class.

Classes

Specifies one or more privilege classes to grant to an administrator. This parameter is required, except when you specify the STGPOOLS parameter. You can specify more than one privilege class by separating each with a comma. Possible classes are:

SYstem

Specifies that you want to grant system privilege to an administrator. A system administrator has the highest level of authority in IBM Spectrum Protect™. A system administrator can issue any administrative command and has authority to manage all policy domains and all storage pools. Do not specify additional privilege classes or the DOMAINS or STGPOOLS parameters when granting system privilege to an administrator. Only a system administrator can grant authority to other administrators.

Policy

Specifies that you want to grant policy privilege to an administrator. If you do not specify the DOMAINS parameter, unrestricted policy privilege is granted. An unrestricted policy administrator can issue commands that affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains. Use the GRANT AUTHORITY command with CLASSES=POLICY and no DOMAINS parameter to upgrade a restricted policy administrator to an unrestricted policy administrator.

STorage

Specifies that you want to grant storage privilege to an administrator. If the STGPOOLS parameter is not specified, unrestricted storage privilege is granted. An unrestricted storage administrator can issue all commands that allocate and control storage resources for the server. An unrestricted storage administrator can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. An unrestricted storage administrator cannot define or delete storage pools. Using the GRANT AUTHORITY command with CLASSES=STORAGE and no STGPOOLS parameter upgrades a restricted storage administrator to an unrestricted storage administrator.

Operator

Specifies that you want to grant operator privilege to an administrator. An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Node

Specifies that you want to grant a node privilege to a user. A user with client node privilege can remotely access a web backup-archive client with an administrative user ID and password if they have been given owner authority or access authority. Access authority is the default for a node privilege class.

Attention: When you specify the node privilege class, you must also specify either the DOMAIN parameter or the NODE parameter, but not both.

AUTHority

Specifies the authority level of a user with node privilege. This parameter is optional.

If an administrator already has system or policy privilege to the policy domain to which the node belongs, this command will not change the administrator's privilege.

Possible authority levels are:

Access

Specifies that you want to grant client access authority to a user with the node privilege class. This is the default when CLASSES=NODE is specified. A user with client access authority can access a web backup-archive client and perform backup and restore actions on that client.

Attention: A user with client access authority cannot access that client from another system by using the -NODENAME or -VIRTUALNODENAME parameter.

A client node can set the REVOKEREMOTEACCESS option to restrict a user that has node privilege with client access authority from accessing a client workstation that is running a web client. This option does not apply to administrators with client owner authority, system privilege, or policy privilege to the policy domain to which the node belongs.

Owner

Specifies that you want to grant client owner authority to a user with the node privilege class. A user with client owner authority can access a web backup-archive client through the web client interface and also access their data from another client using the -NODENAME or -VIRTUALNODENAME parameter.

DOmains

Specifies that you want to grant to the administrator client access or client owner authority to all clients in the specified policy domain. You cannot use this parameter together with the NODE parameter.

NODe

Specifies that you want to grant the administrator client access or client owner authority to the node. You cannot use this parameter together with the DOMAIN parameter.

DOmains

When used with CLASSES=POLICY, specifies that you want to grant restricted policy privilege to an administrator.

Restricted policy privilege permits an administrator to issue a subset of the policy commands for the domains to which the administrator is authorized. You can use this parameter to grant additional policy domain authority to a restricted policy administrator. This parameter is optional. You can specify more than one policy domain by delimiting each policy domain name with a comma.

You can use wildcard characters to specify a name. Authority for all matching policy domains is granted.

STGpools

Specifies that you want to grant restricted storage privilege to an administrator. If the STGPOOLS parameter is specified, then CLASSES=STORAGE is optional.

Restricted storage privilege permits you to issue a subset of the storage commands for the storage pools to which the administrator is authorized. You can use this parameter to grant additional storage pool authority to a restricted storage administrator. This parameter is optional. You can specify more than one storage pool by delimiting each storage pool name with a comma.

You can use wildcard characters to specify a name. Authority for all matching storage pools is granted.

Example: Grant system privilege to an administrator

Grant system privilege to administrator Larry.

```
grant authority larry classes=system
```

Example: Grant access to additional policy domains

Specify additional policy domains that the restricted policy administrator CLAUDIA can manage.

```
grant authority claudia domains=employee_records,progl
```

Example: Provide an administrator with unrestricted storage privilege and restricted policy privilege

Provide administrator TOM with unrestricted storage privilege and restricted policy privilege for the domains whose names start with EMP.

```
grant authority tom classes=storage domains=emp*
```

Example: Grant an administrator authority restricted to a specific node

Grant node privilege to user HELP so that help desk personnel can assist the client node LABCLIENT in backing up or restoring data without having other higher-level IBM Spectrum Protect privileges.

```
grant authority help classes=node node=labclient
```

Related commands

Table 1. Commands related to GRANT AUTHORITY

Command	Description
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.

GRANT PROXYNODE (Grant proxy authority to a client node)

Use this command to grant proxy authority to a client node on the IBM Spectrum Protect™ server.

Target client nodes own the data and agent nodes act on behalf of the target nodes. When granted proxy authority to a target client node, an agent node can perform backup and restore operations for the target node. Data that the agent node stores on behalf of the target node is stored under the target node's name in server storage.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege

Syntax

```
>>-GRant PROXynode TArget--==--target_node_name----->>  
>>-AGent--===--agent_node_name----->><<
```

Parameters

TArget (Required)

Specifies the name of the node that owns the data. Wildcard names cannot be used to specify the target node name.

AGent (Required)

Specifies the name of the node performing operations for the target node. The agent node does not have to be in the same domain as the target node. Wildcard characters and comma-separated lists of node names are allowed.

Example: Grant proxy authority to a client node

Assume that MOE and JOE are agent nodes in a NAS cluster and are used to backup and restore shared NAS data. To create a proxy authority relationship for target node NASCLUSTER, issue the following command:

```
grant proxynode target=nascluster agent=moe,joe
```

Issue the following command on agent node MOE to back up NAS cluster data stored on the E: drive. The name of the target node is NASCLUSTER.

```
dsmc -asnode=nascluster incremental e:
```

Related commands

Table 1. Commands related to GRANT PROXYNODE

Command	Description
QUERY PROXYNODE	Display nodes with authority to act as proxy nodes.
REVOKE PROXYNODE	Revoke proxy authority from an agent node.

HALT (Shut down the server)

Use this command to shut down the server. The HALT command forces an abrupt shutdown, which cancels all the administrative and client node sessions even if they are not completed.

Any transactions in progress interrupted by the HALT command are rolled back when you restart the server. Use the HALT command only after the administrative and client node sessions are completed or canceled. To shut down the server without severely impacting administrative and client node sessions, perform the following steps:

1. Use the DISABLE SESSIONS command to prevent starting new client node sessions.
2. Use the QUERY SESSIONS command to identify any existing administrative and client node sessions.

3. Notify any existing administrative and client node sessions that you plan to shut down the server (you must do this outside of IBM Spectrum Protect™).
4. Use the CANCEL SESSIONS command to cancel any existing administrative or client node sessions.
5. Issue the HALT command to shut down the server and stop any administrative and client node sessions.

Tip:

The HALT command can be replicated using the ALIASHALT server option. Use the server option to define a term other than HALT that performs the same function. The HALT command retains its normal function however, the server option provides an additional method for issuing the HALT command. See ALIASHALT for additional information.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

```
>>-HALT-----<<
```

Parameters

None.

Example: Shut down the server

Shut down the server, either from the server console or from an administrative client. All user activity stops immediately and no new activity can start.

```
halt
```

Related commands

Table 1. Commands related to HALT

Command	Description
CANCEL PROCESS	Cancels a background server process.
CANCEL SESSION	Cancels active sessions with the server.
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY PROCESS	Displays information about background processes.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.

HELP (Get help on commands and error messages)

Use this command to display administrative commands and error messages. You can issue the command from an administrative command line client.

Privilege class

Any administrator can issue this command.

Syntax

3.44 REMOVE commands

Use the REMOVE commands to remove an object.

The following is a list of REMOVE commands:

- * 3.44.1, "REMOVE ADMIN (Delete an administrator)"
- * 3.44.2, "REMOVE NODE (Delete a node or an associated machine node)"

Example: Display help for a specific error message

Display help information about the error message ANR2535E.

```
help 2535
```

```
ANR2535E Command: The node node name cannot be removed or renamed
because it has an associated data mover.
```

```
Explanation: You attempted to remove or rename a node that has an
associated data mover.
```

```
System action: The server does not remove or rename the node.
```

```
User response: To remove or rename the node, delete the associated data
mover and reissue the command.
```

Example: Display help for a specific option

Display the description, syntax, and an example for the COMMETHOD server option.

```
help commethod
```

Example: Display help for a specific utility

Display the description, syntax, and an example for the DSMSERV utility.

```
help dsmserv
```

IDENTIFY DUPLICATES (Identify duplicate data in a storage pool)

Use this command to start or stop processes that identify duplicate data in a storage pool. You can specify the number of duplicate-identification processes and their duration.

When you create a new storage pool for data deduplication, you can specify 0 - 50 duplicate-identification processes. IBM Spectrum Protect™ starts the specified number of duplicate-identification processes automatically when the server is started. If you do not stop them, they run indefinitely.

This command affects only server-side deduplication processing. In client-side data deduplication processing, duplicates are identified on the backup-archive client.

With the IDENTIFY DUPLICATES command, you can start more processes, stop some or all of the processes, and specify an amount of time that the change remains in effect. If you increased or decreased the number of duplicate-identification processes, you can use the IDENTIFY DUPLICATES command to reset the number of processes to the number that is specified in the storage pool definition.

If you did not specify any duplicate-identification processes in the storage pool definition, you can use the IDENTIFY DUPLICATES command to start and stop all processes manually.

This command starts or stops a background process or processes that you can cancel with the CANCEL PROCESS command. To display information about background processes, use the QUERY PROCESS command.

Important:

- You can also change the number of duplicate-identification processes by updating the storage pool definition by using the UPDATE STGPOOL command. However, when you update a storage pool definition, you cannot specify a duration. The processes that you specify in the storage pool definition run indefinitely, or until you issue the IDENTIFY DUPLICATES command, update the storage pool definition again, or cancel a process.

Issuing the IDENTIFY DUPLICATES does not change the setting for the number of duplicate-identification processes in the storage pool definition.

- Duplicate-identification processes can be either active or idle. Processes that are deduplicating files are active. Processes that are waiting for files to deduplicate are idle. Processes remain idle until volumes with data to be deduplicated become

available. Processes stop only when canceled or when you change the number of duplicate-identification processes for the storage pool to a value less than what is specified. Before a duplicate-identification process stops, it must finish the file that it is deduplicating.

The output of the QUERY PROCESS command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files that are processed is nine.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-IDentify DUPLICates--stgpool_name----->
>--+-----+--+----->>
  '-NUMPRocess----number-' '-DURation----minutes-'
```

Parameters

stgpool_name (Required)

Specifies the storage pool name in which duplicate data is to be identified. You can use wildcards.

NUMPRocess

Specifies the number of duplicate-identification processes to run after the command completes. You can specify 0 - 50 processes. The value that you specify for this parameter overrides the value that you specified in the storage pool definition or the most recent value that was specified when you last issued this command. If you specify zero, all duplicate-identification processes stop.

This parameter is optional. If you do not specify a value, the server starts or stops duplicate-identification processes so that the number of processes is the same as the number that is specified in the storage pool definition.

For example, suppose that you define a new storage pool and specify two duplicate-identification processes. Later, you issue the IDENTIFY DUPLICATES command to increase the number of processes to four. When you issue the IDENTIFY DUPLICATES command again without specifying a value for the NUMPROCESS parameter, the server stops two duplicate-identification processes.

If you specified 0 processes when you defined the storage pool definition and you issue IDENTIFY DUPLICATES without specifying a value for NUMPROCESS, any running duplicate-identification processes stop, and the server does not start any new processes.

Remember: When you issue IDENTIFY DUPLICATES without specifying a value for NUMPROCESS, the DURATION parameter is not available. Duplicate-identification processes specified in the storage pool definition run indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.

When the server stops a duplicate-identification process, the process completes the current physical file and then stops. As a result, it might take several minutes to reach the number of duplicate-identification processes that you specified as a value for this parameter.

DURation

Specifies the maximum number of minutes (1 - 9999) that this command remains in effect. At the end of the specified time, the server starts or stops duplicate-identification processes so that the number of processes is the same as the number that is specified in the storage pool definition.

This parameter is optional. If you do not specify a value, the processes that are running after the command is issued run indefinitely. They end only if you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.

For example, if you define a storage pool with two duplicate-identification processes and you issue the IDENTIFY DUPLICATES command with DURATION=60 and NUMPROCESS=4, the server starts two more duplicate-identification processes that run for 60 minutes. At the end of that time, two processes finish the files that they are working on and stop. The two processes that stop might not be the same two processes that started as a result of issuing this command.

The server stops idle processes first. If after stopping all idle processes, more processes need to be stopped, the server notifies active processes to stop.

When the server stops a duplicate-identification process, the process completes the current physical file and then stops. As a result, it might take several minutes to reach the amount of time that you specified as a value for this parameter.

Example: Controlling the number and duration of duplicate-identification processes

In this example, you specified three duplicate-identification processes in the storage pool definition. You use the IDENTIFY DUPLICATES command to change the number of processes and to specify the amount of time the change is to remain in effect.

Table 1. Controlling duplicate-identification processes manually

The storage pool definition specifies three duplicate-identification processes. Using the IDENTIFY DUPLICATES command, you specify...	...and a duration of...	The result is...
2 duplicate-identification processes	None specified	One duplicate-identification process finishes the file that it is working on, if any, and then stops. Two processes run indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	One duplicate-identification process finishes the file that it is working on, if any, and then stops. After 60 minutes, the server starts one process so that three are running.
4 duplicate-identification processes	None specified	The server starts one duplicate-identification process. Four processes run indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	The server starts one duplicate-identification process. At the end of 60 minutes, one process finishes the file that it is working on, if any, and then stops. The additional process started by this command might not be the one that stops when the duration has expired.
0 duplicate-identification processes	None specified	All duplicate-identification processes finish the files that they are working on, if any, and stop. This change lasts indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	All duplicate-identification processes finish the files that they are working on, if any, and stop. At the end of 60 minutes, the server starts three processes.
None specified	Not available	The number of duplicate-identification processes resets to the number of processes that are specified in the storage pool definition. This change lasts indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.

Example: Identify duplicates in a storage pool

Identify duplicates in a storage pool, STGPOOLA, using three duplicate-identification processes. Specify that this change is to remain in effect for 60 minutes.

```
identify duplicates stgpoola duration=60 numprocess=3
```

Related commands

Table 2. Commands related to IDENTIFY DUPLICATES

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.

IMPORT commands

Use the IMPORT commands to import information from export media to an IBM Spectrum Protect™ server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the IMPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

- IMPORT ADMIN (Import administrator information)
- IMPORT NODE (Import client node information)
- IMPORT POLICY (Import policy information)
- IMPORT SERVER (Import server information)

IMPORT ADMIN (Import administrator information)

Use this command to import administrator and authority definitions for one or more administrators from export media to the IBM Spectrum Protect™ server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the IMPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

You can use the QUERY ACTLOG command to view the status of the import operation.

You can also view this information from the server console.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT ADMIN background process is canceled, some of the data is already imported. To display information about background processes, use the QUERY PROCESS command.

Restriction:

- If target and source server levels are not compatible, the operation might not work.
- If the administrator definition that is being imported includes analyst authority, the administrator definition is imported but not the analyst authority. Analyst authority is not valid for servers at V6.1 or later.
- Importing data from a CENTERA device class is not supported. However, files that are being imported can be stored on a CENTERA storage device.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-Import Admin-.*-----.-Preview----No-----
| .,-----| '-Preview----+No---+'
| V         | | '-Yes-'
'---admin_name--'

>--DEVclass----device_class_name----->
      .,-----
      V         |

>--VOLumentname---+---volume_name-+-+----->
      '-FILE:--file_name-'

.-Replacedefs----No-----
>+-----+-----><
  '-Replacedefs----+No--+-'
    '-Yes-'

```

Parameters

admin_name

Specifies the administrators for which you want to import information. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

Preview

Specifies whether you want to preview the results of the import operation, without importing administrator information. This parameter is optional. The following parameter values are supported:

No

Specifies that the information is to be imported.

Yes

Specifies that the operation is previewed but not completed. Information about the number and types of objects that are imported, together with the number of bytes transferred, are reported to the server console and the activity log.

The default value is NO. If you specify YES for the value, you must mount the export volumes.

DEVclass (Required)

Specifies the device class from which import data is to be read.

You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the import runs, IBM Spectrum Protect cancels lower priority operations, such as reclamation, to make a drive available.

VOLumentname (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The following parameter values are supported:

volume_name

Specifies the volume name. To specify multiple volumes, separate names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes that are used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when you specify volumes that are associated with the following device types:

For this device	Specify
Tape	1 - 6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: <div style="display: flex; justify-content: space-around; margin-top: 5px;"> AIX Linux /imdata/mt1. </div> <div style="display: flex; justify-content: space-around; margin-top: 5px;"> Windows d:\program files\tivoli\tsm\data1.dsm. </div>

For this device	Specify
AIX Linux Windows REMOVABLEFILE	AIX Linux Windows 1 - 6 alphanumeric characters.
SERVER	1 - 250 alphanumeric characters.

Replacedefs

Specifies whether to replace administrator definitions on the target server. The following parameter values are supported:

No

Specifies that definitions are not to be replaced.

Yes

Specifies that definitions are to be replaced.

The default value is NO.

Example: Import administrator information from specific tape volumes

From the server, import the information for all defined administrators from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class. Issue the command:

```
import admin devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Import administrator information from tape volumes listed in a file

From the server, import the information for all defined administrators from tape volumes that are listed in the following file:

- AIX** | **Linux** TAPEVOL
- Windows** TAPEVOL.DATA

This file contains these lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that these tape volumes be read by a device that is assigned to the MENU1 device class. Issue the command:

```
AIX | Linux
import admin devclass=menu1 volumenames=file:tapevol

Windows
import admin devclass=menu1 volumenames=file:tapevol.data
```

Related commands

Table 1. Commands related to IMPORT ADMIN

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT POLICY	Restores policy information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT NODE (Import client node information)

Use this command to import client node definitions from a server or sequential media to a target IBM Spectrum Protect™ server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the IMPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

If you specify a domain on the source server and if that policy domain also exists on the target server, the imported nodes get associated with that same policy domain on the target server. Otherwise, imported nodes are associated with the STANDARD policy domain on the target server.

IBM Spectrum Protect servers with retention protection enabled do not allow import operations.

Restrictions:

1. If target and source server levels are not compatible, the operation might not work.
2. Importing data from a CENTERA device class is not supported. However, files that are being imported can be stored on a CENTERA storage device.
3. If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP passwords. Data that is imported from a node that authenticates with an LDAP directory server is inaccessible if the target server is not properly configured. If your target server is not configured, imported data from an LDAP node can still go there. But the target server must be configured to use LDAP in order for you to access the imported data.
4. If target and source server levels are not compatible, the operation might not work.
5. You cannot use a CENTERA device class as the target medium for an export command, or as the source medium for an import command.
6. Incrementally exporting/importing the following types of client data to another IBM Spectrum Protect server is not supported:
 - o VMWare backups where full plus incremental backups need to be periodically, incrementally transferred to another server.
 - o Backups groups where full plus differential backups need to be periodically, incrementally transferred to another server.
 - o **Windows** Windows System State data that is periodically, incrementally transferred to another server.

Full export/import of this data to a new file system on the target is supported by exporting the entire filespace that contains the data. In other words, the export must not use the *FILEDATA=ALLACTIVE*, *FROMDATE*, *TODATE*, or *MERGEFILESACES* options.

The best practice for incrementally transferring this type of data between two servers is to use Node Replication.

You can use the QUERY ACTLOG command to view the status of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT NODE background process is canceled, some of the data might already be imported. To display information about background processes, use the QUERY PROCESS command.

For a server that has clients with support for Unicode, you can get the server to convert the file space name that you enter, or use the following parameters:

- HEXFILESACE
- UNIFILESACE

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

.*-----
>>-Import Node----->
| .-,-----|
| V          |
|'---node_name-+-'
>----->
| .-,-----|
| V          |
|'-FILESpace-----file_space_name-+-'
>----->
| .-,-----|
| V          |
|'-HEXFILESpace-----file_space_name-+-'
>----->
| .-,-----|
| V          |
|'-UNIFILESpace-----file_space_name-+-'
>----->
| .-,-----|
| V          |
|'-DObains-----domain_name-+-'
.-FILEData-----None-----.-Preview-----No-----
>----->
|'-FILEData-----+All-----+|'-Preview-----+No-+-|
|+None-----+|'+Yes-|
|+ARchive-----+
|+Backup-----+
|+BACKUPActive-+
|+ALLActive-----+
|'-SPacemanaged-'
.-Dates-----Absolute-----
>>-DEVclass-----device_class_name----->
|'-Dates-----+Absolute-+-|
|'+Relative-|
.-,-----
V          |
>>-VOLumenames-----+---volume_name-+-+----->
|'-FILE:--file_name-'
.-Replacedefs-----No-----
>----->
|'-Replacedefs-----+No-+-|
|'+Yes-|
.-MERGEfilespaces-----No-----
>----->
|'-MERGEfilespaces-----+No-+-|
|'+Yes-|
.-PROXynodeassoc-----No-----
>----->
|'-PROXynodeassoc-----+No-+-|
|'+Yes-|

```

Parameters

node_name

Specifies the client nodes for which you want to import information. This parameter is optional.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. All matching nodes are included in the list.

FILESpace

Specifies file space names for which you want to import information. This parameter is optional. The default is all file spaces.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

Important:

1. Existing file spaces are not replaced. New file spaces are created when identical names are encountered. However, this new name might match an existing name on the client node, which can have file spaces that are not yet backed up to the server.
2. This parameter is only specified for non-Unicode file spaces. To import all file spaces that are both Unicode and non-Unicode, use the FILEDATA=ALL parameter without the FILESPACE and UNIFILESPACE parameters.

DOmains

Specifies the policy domains from which to import node information. These domains must be included in the data that was exported. This parameter is optional. The default is all domains that were exported.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

FILEData

Specifies the type of files that can be imported for all nodes that are specified and found on the export media. This parameter is optional. The default value is NONE.

If you are importing from sequential media, the device class that is used by the file data is determined by the device class for the storage pool. If it is the same device class that is specified in this command, two drives are needed to import the node information. The mount limit for the device class must be at least 2.

The following descriptions mention *active* and *inactive* backup file copies. An active backup file copy is the most recent backup copy for a file that still exists on the client workstation. All other backup file copies are called inactive copies. The parameter supports the following values:

ALL

The server imports all backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client. The file spaces that are included are both Unicode and non-Unicode.

None

Only node definitions are imported. The server does not import any files.

ARchive

The server imports only archived files.

Backup

The server imports only backup versions, whether active or inactive.

BACKUPActive

The server imports only active backup versions. These active backup versions are the active versions in the IBM Spectrum Protect database at the time that the IMPORT command is issued.

ALLActive

The server imports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client. The active backup versions are the active versions in the IBM Spectrum Protect database at the time that the IMPORT command is issued.

SPacemanaged

The server imports only files that were migrated by an IBM Spectrum Protect for Space Management client.

Preview

Specifies whether to preview the results of the import operation, without importing information. The PREVIEW=YES option requires that you mount the export volumes. The following values are supported:

No

Specifies that the node information is to be imported.

Yes

Specifies that you want to preview the results of the import operation, without importing files. Information is reported to the server console and the activity log.

This parameter is optional. The default value is NO.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the import runs, the server cancels lower priority operations, such as identify duplicates, to make a drive available.

Dates

Specifies whether the dates for the file copies are set as the same date when the files were exported, or is adjusted to the import date.

This parameter supports the following values:

Absolute

The dates for file copies are set to the values specified when the files were exported.

Relative

The dates for file copies are adjusted to the import date.

The default value is ABSOLUTE.

If the export media is idle for some time after export, for example; if it is sitting on a shelf for six months, the original backup, or archive dates might be old enough to trigger the file copies to expire immediately when the data is imported into a server. The RELATIVE specification for this value adjusts for time that is elapsed since export so that the file copies are not immediately expired.

For example, assume that an export tape contains an archive file copy that was archived five days before the export operation. If the media is saved for six months and then imported, the archive file look like it is inserted six months and five days ago by default, the (DATES=ABSOLUTE) and might expire immediately depending on the retention value that is specified in the file's management class. Specifying DATES=RELATIVE results in resetting the archive date for the file to five days ago during import. The DATES=RELATIVE parameter thus adjusts file backup and archive dates for the time that elapsed since the export operation occurred.

VOLUMenames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The parameter supports the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes that are used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when you specify volumes that are associated with the following device types:

For this device	Specify
Tape	1 - 6 alphanumeric characters.
FILE	<p>AIX Linux Any fully qualified file name string. An example is /imdata/mt1.</p> <p>Windows Any fully qualified file name string. For example, d:\program files\tivoli\tsm\data1.dsm.</p>
AIX Linux Windows REMOVABLEFILE	AIX Linux Windows 1 - 6 alphanumeric characters.
SERVER	1 - 250 alphanumeric characters.

Replacedefs

Specifies whether to replace definitions on the target server. The default value is NO. The parameter supports the following values:

No

Objects are not to be replaced.

Yes

Objects are to be replaced.

HEXFILESspace

Specifies the hexadecimal representation of the file space names in UTF-8 format. Separate multiple names with commas and no intervening spaces. This parameter is optional.

To view the hexadecimal representation of a file space name, you can use the QUERY FILESPACE command with FORMAT=DETAILED.

UNIFILESpace

Specifies that the file spaces that are known to the server are Unicode enabled. The server converts the names that you enter from the server code page to the UTF-8 code page to find the file spaces to import. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. A wildcard character can be used to specify a name. This parameter is optional.

MERGEfilespace

Specifies whether IBM Spectrum Protect merges client files into existing file spaces on the target server (if they exist), or if IBM Spectrum Protect generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No

Specifies that IBM Spectrum Protect generates a new file space name for imported data on the target server if file spaces with the same name exists.

PROXynodeassoc

Specifies whether proxy node associations are imported. This parameter is optional. The default value is NO.

Example: Import client node information from tapes

From the server, import client node information from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class.

```
import node devclass=menu1 volumenames=tape01,tape02,tape03
```

Example: Import client node information from tapes listed in a file

AIX | **Linux** From the server, import client node information from tape volumes that are listed in a file named TAPEVOL.

Windows From the server, import client node information from tape volumes that are listed in a file named TAPEVOL.DATA.

This file contains these lines:

```
TAPE01  
TAPE02  
TAPE03
```

Specify that these tape volumes be read by a device that is assigned to the MENU1 device class. **AIX** | **Linux**

```
import node devclass=menu1 volumenames=file:tapevol
```

Windows

```
import node devclass=menu1 volumenames=file:tapevol.data
```

Example: Import the active backup for a client node

From the server, import the active backup versions of file data for client node JOE from tape volume TAPE01. The file space is Unicode.

```
import node joe unifiespace=\\joe\c$ filedata=backupactive devclass=menu1  
volumenames=tape01
```

Related commands

Table 1. Commands related to IMPORT NODE

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.

Command	Description
EXPORT NODE	Copies client node information to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
IMPORT POLICY	Restores policy information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT POLICY (Import policy information)

Use this command to import policy domain information from sequential export media to the IBM Spectrum Protect™ server. IBM Spectrum Protect servers with retention protection enabled do not allow import operations.

IBM Spectrum Protect client data can be moved between servers with export and import processing, if the same removable media type is supported on both platforms.

Restriction:

1. If target and source server levels are not compatible, the import operation might not work.
2. Importing data from a CENTERA device class is not supported. However, files that are imported can be stored on a CENTERA storage device.

You can use the QUERY ACTLOG command to view the status of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT POLICY background process is canceled, some of the data is already imported. To display information about background processes, use the QUERY PROCESS command.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-Import Policy-+-----+----->
      .-*-----*
      | .-,-----, |
      | v             | |
      |---domain_name---|
      .-Preview----No-----
>>-+-----+---DEVclass----device_class_name----->
      '-Preview----+No---+'
          '-Yes-'

      .-,-----,
      v             |
>>-VOLumenames----+---volume_name-+----->
          '-FILE:--file_name-'

      .-Replacedefs----No-----

```

```
>-----<
'-Replacedefs-----+--No--+-'
'-Yes-'
```

Parameters

domain_name

Specifies the policy domains for which information is to be imported. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. The default (*) is all policy.

Preview

Specifies whether you want to preview the results of the import operation without importing information. This parameter supports the following values:

No

Specifies that the information is to be imported.

Yes

Specifies that the operation is previewed but not completed. Information is reported to the server console and the activity log.

The PREVIEW=YES option requires that you mount the export volumes. This parameter is optional. The default value is NO.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the import runs, IBM Spectrum Protect cancels lower priority operations, such as reclamation, to make a drive available.

VOLumentnames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. This parameter supports the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when you specify volumes that are associated with the following device types:

For this device	Specify
Tape	1 - 6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: <ul style="list-style-type: none"> AIX Linux /imdata/mt1 Windows d:\program files\tivoli\tsm\data1.dsm.
AIX Linux Windows REMOVABLEFILE	AIX Linux Windows 1 - 6 alphanumeric characters.
SERVER	1 - 250 alphanumeric characters.

Replacedefs

Specifies whether to replace policy definitions on the target server. This parameter supports the following values:

Yes

Specifies that objects are to be replaced by the imported objects.

No

Specifies that objects are not to be replaced by imported objects.

The default value is NO.

Example: Import policy information from specific tape volumes

From the server, import the information for all defined policies from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class.

```
import policy devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Import policy information from tape volumes listed in a file

From the server, import the information for all defined policies from tape volumes that are listed in a file that is named thus:

- **AIX** | **Linux** TAPEVOL
- TAPEVOL.DATA

Specify that these tape volumes be read by a device that is assigned to the MENU1 device class. The file contains the following lines:

```
TAPE01
TAPE02
TAPE03
```

AIX | **Linux**

```
import policy devclass=menu1 volumenames=file:tapevol
```

Windows

```
import policy devclass=menu1 volumenames=file:tapevol.data
```

Related commands

Table 1. Commands related to IMPORT POLICY

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT POLICY	Copies policy information to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT SERVER (Import server information)

Use this command to copy all or part of the server control information and specified client file data from export media to the IBM Spectrum Protect™ server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The IBM Spectrum Protect server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the IMPORT command, set the password by issuing the UPDATE ADMIN or UPDATE NODE command.

IBM Spectrum Protect servers with retention protection enabled do not allow import operations.

Restrictions:

- If target and source server levels are not compatible, the operation might not work.
- Importing data from a CENTERA device class is not supported. However, files that are imported can be stored on a CENTERA storage device.
- If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP passwords. Server data that is exported from a node that authenticates with an LDAP directory server is inaccessible if the target server

is not properly configured. If your target server is not configured, exported data from an LDAP node can still go there. But the target server must be configured to use LDAP in order for you to access the data.

- Incrementally exporting or importing the following types of client data to another IBM Spectrum Protect server is not supported:
 - VMware backups where full plus incremental backups need to be periodically, incrementally transferred to another server
 - Backups groups where full plus differential backups must be periodically, incrementally transferred to another server
 - Windows System State data that is periodically, incrementally transferred to another server

Full export or import of this data to a new file system on the target is supported by exporting the entire file space that contains the data. The export must not use the FILEDATA=ALLACTIVE, FROMDATE, TODATE, or MERGEFILESPPACES parameters.

Using node replication to incrementally transfer this type of client data between two servers is optimal.

You can also initiate an import of server information and client file data directly from the originating server. For more information, see the EXPORT commands.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT SERVER background process is canceled, some of the data is already imported. To display information about background processes, use the QUERY PROCESS command.

Limitation: The IBM Spectrum Protect server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue after the import or node replication operation, update the fields with the appropriate UPDATE commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-FILEData-----None----- .
>>-Import Server-----+-----+----->
      '-FILEData-----+Al-----'
                          +-None-----+
                          +-ARchive-----+
                          +-Backup-----+
                          +-BACKUPActive+
                          +-ALLActive-----+
                          '-SPacemanged-'

      .-Preview-----No----- .
>>-+-----+-----DEVclass-----device_class_name----->
      '-Preview-----+No--+-'
                          '-Yes-'

      .-Dates-----Absolute----- .
>>-+-----+-----+----->
      '-Dates-----+Absolute+-'
                          '-Relative-'

      .-,----- .
      V          |
>>-VOLumenames-----+---volume_name+---+----->
                          '-FILE:--file_name-'

      .-Replacedefs-----No----- .
>>-+-----+-----+----->
      '-Replacedefs-----+No--+-'
                          '-Yes-'

```



```

.-MERGEfilespace-----No----- .
>-----+----->
'-MERGEfilespace-----+No--+-'
      '-Yes-'

.-PROXynodeassoc-----No----- .
>-----+-----><
'-PROXynodeassoc-----+No--+-'
      '-Yes-'

```

Parameters

FILEData

Specifies the type of files that can be imported for all nodes that are defined to the server. This parameter is optional. The default value is NONE.

The device class that is used to access the file data is determined by the device class for the storage pool. If it is the same device class that is specified in this command, two drives are needed to import information. The mount limit for the device class must be set to at least 2.

The following descriptions mention active and inactive backup file copies. An active backup file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. This parameter supports the following values:

ALL

IBM Spectrum Protect imports all backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client.

None

IBM Spectrum Protect does not import files, only node definitions.

ARchive

IBM Spectrum Protect imports only archived files.

Backup

IBM Spectrum Protect imports only backup versions, whether the versions are active or inactive.

BACKUPActive

IBM Spectrum Protect imports only active backup versions. These active backup versions are the active versions in the IBM Spectrum Protect database at the time that the IMPORT command is issued.

ALLActive

IBM Spectrum Protect imports all active backup versions of files, all archived files, and all files that were migrated by an IBM Spectrum Protect for Space Management client. The active backup versions are the active versions in the IBM Spectrum Protect database at the time that the IMPORT command is issued.

SPacemanaged

IBM Spectrum Protect imports only files that were migrated by an IBM Spectrum Protect for Space Management client.

Preview

Specifies whether to preview the results of the import operation, without importing information. This parameter supports the following values:

No

Specifies that the server information is to be imported.

Yes

Specifies that the operation is previewed but not completed. Information is transferred to the server console and the activity log.

This parameter is optional. The default value is NO. If the PREVIEW=YES option is specified, you must mount the export volumes.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the DISK, NAS, or CENTERA device classes.

If all drives for the device class are busy when the import runs, IBM Spectrum Protect cancels lower priority operations, such as reclamation, to make a drive available.

Dates

Specifies whether the dates for the file copies are set as the same date when the files were exported, or is adjusted to the import date.

If the import media is idle for some time after export, for example; if it is sitting on a shelf for six months, the original backup, or archive dates might be old enough to trigger the file copies to expire immediately when the data is imported into a server. The RELATIVE specification for this value adjusts for time that is elapsed since export so that the file copies are not immediately expired.

For example, assume that an import tape contains an archive file copy that was archived five days before the export operation. If the export media are saved for six months and then imported, the archive file looks like it is inserted six months and five days ago by default (DATES=ABSOLUTE) and might expire immediately depending upon the retention value that is specified in the file's management class. Specifying DATES=RELATIVE results in resetting the archive date for the file to five days ago during import. DATES=RELATIVE parameter thus adjusts file backup and archive dates for the time that elapsed since the export operation occurred.

This parameter supports the following values:

Absolute

The dates for file copies are set to the values specified when the files were exported.

Relative

The date for file copies are adjusted to the date of import.

The default value is ABSOLUTE.

VOLumentnames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. This parameter supports the following values:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:file_name

Specifies the name of a file that contains a list of volumes that are used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when you specify volumes that are associated with the following device types:

For this device	Specify
Tape	1 - 6 alphanumeric characters.
FILE	<p>AIX Linux Any fully qualified volume or file name string. An example is /imdata/mt1.</p> <p>Windows Any fully qualified volume or file name string. For example, d:\program files\tivoli\tsm\data1.dsm.</p>
AIX Linux Windows REMOVABLEFILE	AIX Linux Windows 1 - 6 alphanumeric characters.
SERVER	1 - 250 alphanumeric characters.

Replacedefs

Specifies whether to replace objects on the server. Existing file spaces are not replaced. New file spaces are created when identical names are encountered. This parameter supports the following values:

No

Specifies that objects are not to be replaced by imported objects.

Yes

Specifies that objects are to be replaced by the imported objects.

The default value is NO.

MERGEfilespace

Specifies whether IBM Spectrum Protect merges client files into existing file spaces on the target server (if they exist), or if IBM Spectrum Protect generates new file space names. You cannot merge non-Unicode and Unicode file spaces together. This parameter supports the following values:

No

Specifies that IBM Spectrum Protect generates a new file space name for imported data on the target server if file spaces with the same name exist.

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

The default is NO.

PROXynodeassoc

Specifies whether proxy node associations are imported. This parameter is optional. The default value is NO.

Example: Import the information for all defined servers from specific tapes

From the server, import the information for all defined servers from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class.

```
import server devclass=menu1 volumenames=tape01,tape02,tape03
```

AIX

Linux

Example: Import information for all defined servers from specific tapes and specify files are merged into existing file spaces

From the server, import the information for all defined servers from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device that is assigned to the MENU1 device class and that client files be merged into file spaces on the target server if file spaces of the same names exist.

```
import server devclass=menu1 volumenames=tape01,tape02,tape03 mergefilespace=yes
```

Example: Import information for all defined servers from tapes listed in a file

From the server, import the information for all defined servers from tape volumes that are listed in a file named TAPEVOL. Specify that the tape volumes are read by a device that is assigned to the MENU1 device class. The input file contains these lines:

```
TAPE01  
TAPE02  
TAPE03
```

```
import server devclass=menu1 volumenames=file:tapevol
```

Windows

Example: Import information for all defined servers from tapes listed in a file

From the server, import the information for all defined servers from tape volumes that are listed in a file named TAPEVOL.DATA. Specify that the tape volumes are read by a device that is assigned to the MENU1 device class. The input file contains these lines:

```
TAPE01  
TAPE02  
TAPE03
```

```
import server devclass=menu1 volumenames=file:tapevol.data
```

Related commands

Table 1. Commands related to IMPORT SERVER

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
IMPORT NODE	Restores client node information from external media.

Command	Description
IMPORT POLICY	Restores policy information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

INSERT MACHINE (Insert machine characteristics information or recovery instructions)

Use this command to add client machine characteristics or recovery instructions to existing machine information in the database.

You can write a program to read files containing the information and generate the appropriate INSERT MACHINE commands.

You can use QUERY commands to retrieve the information if a disaster occurs.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-INsert MACHine--machine_name--sequence_number----->
>--+CHARacteristics---text-----+-----><
  '-RECOVERYInstructions---text-'
```

Parameters

machine_name (Required)

Specifies the name of the client machine.

sequence_number (Required)

Specifies the sequence number for the line of text in the database.

CHARacteristics

Specifies machine characteristics information. You must specify the characteristics or recovery instructions, but not both. Enclose the text in quotation marks if it contains blank characters. The text can be up to 1024 characters.

RECOVERYInstructions

Specifies recovery instructions. You must specify the characteristics or recovery instructions, but not both. Enclose the text in quotation marks if it contains blank characters. The text can be up to 1024 characters.

Example: Update a machine's information

For the machine DISTRICT5, insert this characteristics text on line 1: "Machine owner is Mary Smith".

```
insert machine district5 1
characteristics="Machine owner is Mary Smith"
```

Related commands

Table 1. Commands related to INSERT MACHINE

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DELETE MACHINE	Deletes a machine.
QUERY MACHINE	Displays information about machines.

Related information:

[Specifying information about your server and client node machines](#)

ISSUE MESSAGE (Issue a message from a server script)

Use this command with return code processing in a script to issue a message from a server script to determine where the problem is with a command in the script.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-ISSUE MESSAGE--message_severity--message_text-----<<
```

Parameters

message_severity (Required)

Specifies the severity of the message. The message severity indicators are:

- I Information. ANR1496I is displayed in the message text.
- W Warning. ANR1497W is displayed in the message text.
- E Error. ANR1498E is displayed in the message text.
- S Severe. ANR1499S is displayed in the message text.

message_text (Required)

Specifies the description of the message.

Example: Issue a message from a server script

Assume you have a script called `backupsript` that quiesces a client's database, takes a backup of that database, and then restarts the client's database. For illustration, your script results in a non-zero return code. Use the `ISSUE MESSAGE` command with the message severity and message text. The following is an example of a server script that calls `backupsript` on the client machine and issues messages based on the return code from `backupsript`.

```
issue message i "Starting backup"
define clientaction nodename action=command objects="c:\backupsript" wait=yes
if (101) goto qfail
if (102) goto qwarn
if (103) goto backupf
if (104) goto restartf
issue message i "Backup of database complete"
exit
qfail: issue message e "Quiesce of database failed"
exit
qwarn: issue message w "Quiesce of database failed, taking fuzzy backup"

exit
backupf: issue message e "Backup of database failed"
exit
restartf: issue message s "Database restart failed"
exit
```

Command

```
issue message e "quiesce of database failed"
```

Related commands

Table 1. Commands related to `ISSUE MESSAGE`

Command	Description
---------	-------------

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

LABEL LIBVOLUME (Label a library volume)

Use this command to label tape volumes or, in an automated library, to label the volumes automatically as they are checked in. With this command, the server uses the full-length label with which the volumes are often pre-labeled.

Restriction: Use this command only for MANUAL, SCSI, ACSLS, and 349X libraries. The command processing does not wait for a drive to become available, even if the drive is only in the IDLE state. If necessary, you can make a library drive available by issuing the DISMOUNT VOLUME command to dismount the volume in that particular drive. When the library drive becomes available, you can reissue the LABEL LIBVOLUME command.

For detailed and current drive and library support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

AIX | **Linux**

To use the LABEL LIBVOLUME command, at least one drive must exist that is not in use by another IBM Spectrum Protect™ process. This includes idle volumes that are mounted. If necessary, use the DISMOUNT VOLUME command to dismount the idle volume to make that drive available.

By default, the LABEL LIBVOLUME command does not overwrite an existing label. However, if you want to overwrite an existing label, you can specify the `OVERWRITE=YES` option.

Attention:

- By overwriting a volume label, you destroy all data on the volume. Use caution when you overwrite volume labels to avoid deleting valid data.
- The labels on VolSafe volumes can be overwritten only once. Therefore, use the LABEL LIBVOLUME command only once for VolSafe volumes. You can guard against overwriting the label by using the `OVERWRITE=NO` option with the LABEL LIBVOLUME command.

When you use the LABEL LIBVOLUME command, you can identify the volumes to be labeled in one of the following ways:

- Explicitly name one volume.
- Enter a range of volumes by using the VOLRANGE parameter.
- Use the VOLLIST parameter to specify a file that contains a list of volume names or to explicitly name one or more volumes.

For automated libraries, you are prompted to insert the volume in the entry/exit slot of the library.

When virtual input/output (VIO) is enabled, volumes that are in the I/O station are no longer in entry/exit ports. To ensure that the volumes can be processed, move them from the I/O station to VIO slots. If no I/O convenience station is available, insert the volume into an empty slot.

For manual libraries, you are prompted to load the volume directly into a drive.

Tip: To automatically label tape volumes, you can use the AUTOLABEL parameter on the DEFINE LIBRARY and UPDATE LIBRARY commands. By using the AUTOLABEL parameter, you eliminate the need to pre-label a set of tapes. This method is more efficient than using the LABEL LIBVOLUME command, which requires you to mount volumes separately. If you use the AUTOLABEL parameter with a SCSI library, you must check in tapes by specifying `CHECKLABEL=BARCODE` on the CHECKIN LIBVOLUME command. The AUTOLABEL parameter defaults to YES for all non-SCSI libraries and to NO for SCSI libraries.

Windows

To label volumes with the LABEL LIBVOLUME command, specify the CHECKIN parameter.

To automatically label tape volumes in SCSI-type libraries, use the AUTOLABEL parameter on the DEFINE LIBRARY and UPDATE LIBRARY commands. By using this parameter, you eliminate the need to pre-label a set of tapes. This method is also more efficient than using the LABEL LIBVOLUME command, which requires you to mount volumes separately. If you use the AUTOLABEL parameter, you must check in tapes by specifying CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

A label cannot include embedded blanks or periods and must be valid when used as a file name on the media.

You must label CD-ROM, Zip, or Jaz volumes with the device utilities from the manufacturer or the Windows utilities. IBM Spectrum Protect does not provide utilities to format or label these media types. The operating system utilities include the Disk Administrator program (a graphical user interface) and the label command.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for a manual library

```
>>-LABEL LIBVolume--library_name-----volume_name----->
      .-OVERWRITE-----No----- .-WAITTime-----60----.
>--+-----+-----+-----+-----><
      '-OVERWRITE-----+No--+-' '-WAITTime-----value-'
                '-Yes-'
```

Syntax for a SCSI library

```
>>-LABEL LIBVolume--library_name----->
>-----+volume_name-----+----->
      '-SEARCH-----+Yes--| A |-----LABELSource-----+Barcode-----+-'
                '-Bulk--| A |-'                +-Prompt-----+
                'Vollist--| B |-'
                .-OVERWRITE-----No-----.
>--+-----+-----+-----+----->
      '-CHECKIN-----+SCRatch+-' '-OVERWRITE-----+No--+-'
                '-PRivate-'                '-Yes-'
      .-WAITTime-----60----.
>--+-----+-----+-----+-----><
      '-WAITTime-----value-'
```

A (SEARCH=Yes, SEARCH=Bulk)

```
|--+VOLRange-----volume_name1,volume_name2--+-----|
|          .-,-----|
|          V          |
'-VOLList-----+volume_name+-----'
                '-FILE:--file_name-'
```

B (LABELSource=Vollist)

```
      .-,-----.
      V          |
|--VOLList-----+volume_name+-----|
                '-FILE:--file_name-'
```

Syntax for a 349X library

```
>>-LABEL LIBVolume--library_name----->
>-----+volume_name-----+----->
      '-SEARCH-----Yes-----| A |---'
                .-OVERWRITE-----No-----.
```


If you do not specify a value for this parameter, the command labels the volume, but does not check it in. If you do not specify a value for this parameter and you want to check in the volume, you must issue the CHECKIN LIBVOLUME command.

SEARCH

Specifies that the server searches the library for usable volumes to label. This parameter applies to SCSI, 349X, and ACSLS libraries.

The following values are valid:

Yes

Specifies that the server labels only volumes that are stored in the library, unless the volume is already labeled or its bar code cannot be read.

If you specify the LABELSOURCE=PROMPT option, the volume is moved into the drive from its location in the library or entry and exit ports. The server prompts you to issue the REPLY command that contains the label string, and that label is written to the tape.

Bulk

Specifies that the server searches the library entry/exit ports for usable volumes to label. This option is only valid for SCSI libraries.

If you specify LABELSOURCE=BARCODE, the volume bar code is read. Then, the tape is moved from its location in the library or in the entry/exit ports to a drive where the bar code label is written. After the tape is labeled, it is moved back to its location in the library, to the entry/exit ports, or to a storage slot if the CHECKIN option is specified. For bar code support to work correctly for libraries that are supported by IBM Spectrum Protect, the IBM Spectrum Protect server and the device driver must be at the same level. Bar code support is available for libraries that are supported by IBM Spectrum Protect and that use the IBM Spectrum Protect device driver or the IBM® Magstar® or LTO Ultrium device driver.

Tip: You can use the VOLRANGE or VOLLIST parameter to limit the search.

VOLRange

Specifies a range of volume names that are separated by a comma. Use this parameter to limit the search for volumes to be labeled when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are within the specified range, the command completes without errors.

You can specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
volrange=bar110,bar130	The 21 volumes are labeled: bar110, bar111, bar112,...bar129, bar130.
volrange=bar11a,bar13a	The 3 volumes are labeled: bar11a, bar12a, bar13a.
volrange=123400,123410	The 11 volumes are labeled: 123400, 123401, ...123409, 123410.

VOLLIST

Specifies a list of volumes. Use this parameter to limit the search for volumes to be labeled when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are in the list, the command completes without errors. The VOLLIST parameter can also be the source of names to be used to label volumes if the LABELSOURCE parameter is set to VOLLIST. If LABELSOURCE=VOLLIST, you must specify the VOLLIST parameter.

The following values are valid:

volume_name

Specifies the names of one or more values that are used for the command. For example: VOLLIST=TAPE01, TAPE02.

FILE:file_name

Specifies the name of a file that contains a list of volumes for the command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volume TAPE01, TAPE02 and TAPE03, create a file that is named TAPEVOL that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can specify the volumes for the command as follows: `VOLLIST=FILE:TAPEVOL`.

Remember: The file name is case-sensitive.

LABELSource

Specifies how or whether the server reads sequential media labels of volumes. This option is only valid for SCSI libraries. Specify this parameter only when `SEARCH=YES` or `SEARCH=BULK`.

You can specify the following values:

Prompt

The server prompts for volume names as necessary.

Barcode

The server attempts to read the bar code label. If the attempt fails, the server does not label the volume and displays a message.

Important: For bar code support to work properly, the appropriate device drivers must be installed for the libraries.

Vollist

This option applies only to SCSI libraries. The server attempts to read the specified file or list of files. If the attempt fails, the server does not label the volumes and displays a message.

OVERWRITE

Specifies whether the server attempts to overwrite existing labels. This parameter is optional. The default is `NO`. You can specify the following values:

No

Specifies that the server labels only unlabeled volumes. For StorageTek VolSafe volumes, the value must be `NO`.

Yes

Specifies that the server overwrites existing labels only if both the existing label and the prompted or bar code label are not already defined in either the server storage pool or volume history list.

WAITTime

Specifies the number of minutes that the server waits for you to reply or respond to a request. Specify a value in the range 0-9999. If you want to be prompted by the server, specify a wait time greater than zero. The default value is 60 minutes. For example, suppose that the server prompts you to insert a tape into the entry/exit port of a library. If you specified a wait time of 60 minutes, the server issues a request and wait 60 minutes for you to reply. Alternatively, suppose that you specify a wait time of 0. If you inserted a tape, a wait time of zero causes the operation to continue without prompting. If you did not insert a tape, a wait time of zero causes the operation to fail.

Example: Automatically label library volumes

Label tapes in a SCSI library named `AUTO` automatically as you are checking in the volumes.

```
label libvolume auto checkin=scratch search=yes labelsource=barcode  
overwrite=yes
```

Example: Label sequential library volumes

Label 3 volumes from `bar11a` to `bar13a` in a SCSI library named `ABC`. When you issue the following command, the three volumes are labeled: `bar11a`, `bar12a`, `bar13a`.

```
label libvolume abc checkin=scratch search=yes volrange=bar11a,bar13a  
labelsource=barcode
```

Related commands

Table 1. Commands related to LABEL LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CANCEL PROCESS	Cancels a background server process.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE LIBRARY	Defines an automated or manual library.

Command	Description
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
UPDATE LIBVOLUME	Changes the status of a storage volume.

LOAD DEFALERTTRIGGERS (Load the default set of alert triggers)

Use this command to load the default set of alert triggers to the IBM Spectrum Protect™ server.

For a newly installed server, a default set of messages is defined to trigger alerts. You can modify or delete default alert triggers. Use this command to complete the following tasks:

- Load the default set of alert triggers, restoring any that were deleted.
- Replace all alert triggers with the original default set.

By default, this command does not delete other alert triggers that were created, and does not replace default alert triggers that were modified. To delete all alert triggers and restore the original set of default alert triggers, specify RESET=yes.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-L0ad DEFALerttriggers--+-REset-----No-----+----->>
                          '-REset-----+No---+'
                          '-Yes-'

```

Parameters

REset

Specifies whether you want to replace all of your alert triggers with the default set of alert triggers. This parameter is optional. The default value is No. Possible values are:

No

Specifies that the default alert triggers are added only. The original default alert triggers are added to the server. Existing triggers are not deleted. If a default trigger exists on the server, it is not replaced or modified.

Yes

Specifies that the alert triggers are restored to the original defaults. All alert triggers are deleted and then the original set of default alert triggers are added.

Example: Load the default alert triggers on the server

Load the default triggers to restore any that were deleted. Issue the command:

```
load defalertriggers
```

Example: Replace all alert triggers on the server with the default alert triggers

Delete all alert triggers on the server and replace them with the original defaults. Issue the command:

```
load defalertriggers reset=yes
```


Example: Lock all nodes that authenticate to the local IBM Spectrum Protect database

Issue the following command to lock all nodes that authenticate with the IBM Spectrum Protect server:

```
lock node * authentication=local
```

Related commands

Table 1. Commands related to LOCK NODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.

LOCK PROFILE (Lock a profile)

Use this command on a configuration manager to temporarily lock a profile so that configuration information is not distributed to subscribing managed servers.

You can use this command when you are making multiple updates to your configuration and do not want to distribute this information until the changes are completed.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-LOCK PROFILE--profile_name--+-60-----+-----><
                               '-minutes-'
```

Parameters

profile_name (Required)

Specifies the profile to lock. You can use wildcard characters to indicate multiple names.

minutes

Specifies the time, in minutes, before IBM Spectrum Protect™ unlocks the configuration profile. Specify an integer from 0 to 10000. The default is 60 minutes. If you specify 0, the configuration profile will not unlock automatically. Use the UNLOCK PROFILE command to unlock the profile before the time period elapses, or to unlock it if you have specified a value of 0.

This parameter is optional.

Example: Lock a profile for a specific amount of time

Lock a profile named DELTA for 30 minutes.

```
lock profile delta 30
```

Related commands

Table 1. Commands related to LOCK PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.

Command	Description
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

MACRO (Invoke a macro)

Use this command to invoke a file from the administrative command line that contains one or more IBM Spectrum Protect™ administrative commands to be performed.

Restriction: Use this command with administrative command-line clients only.

A macro is a file that contains one or more IBM Spectrum Protect administrative commands. You can only issue a macro from the administrative client in batch or interactive mode. A macro is stored as a file on the administrative client machine (or system). Macros are not distributed across servers and cannot be scheduled on the server.

Creating a macro to enter commands can be helpful when you want to issue commands that are used repeatedly, to issue commands that contain several parameters, or to process related commands in a specific order. After you create a macro, you can update the information it contains and use it again, or you can copy the macro file, make changes to the copy, and then run the copy.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-MACRO--macro_name--+-----+----->>
      | .-----+-----+ |
      | v           | |
      |---substitution_value+--'
```

Parameters

macro_name (Required)

Specifies the name of the macro.

substitution_value

Specifies the value for a substitution variable in a macro. When you use a substitution variable, you can reuse a macro whenever you need to perform the same task for different objects or with different parameter values. To specify a value that contains blanks, you must enclose the value in quotation marks. This parameter is optional.

Example: Create a macro to register a new administrator

Create a macro file named REGNG. Use the macro to register and grant authority to a new administrator. Write the macro as follows:

```
/* Register and grant authority to a new administrator */
REGister Admin jones passwd -
CONtactinfo="x1235"
GRant AUTHority jones -
CLasses=Policy
```

Issue the following command to run the macro:

Example: Write a macro using substitution variables

Create a macro file named AUTHRG, containing substitution variables, to register and grant authority to a new administrator. Write the macro as follows:

```
/* Register and grant authority to a new administrator */
REGister Admin %1 %2 - /* Enter userid and password */
CONtact=%3 /* Enter contact info (in quotes if nec.) */
GRant AUTHority %1 - /* Server uses variable already */
- /* defined by you */
CLasses=%4 /* Enter the privilege class */
```

Issue a command similar to the following, entering the values you want to pass to the server to process the command when you run the macro.

```
macro authrg.mac jones passwd x1235 Policy
```

Related commands

Table 1. Commands related to MACRO

Command	Description
COMMIT	Makes changes to the database permanent.
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.

Related concepts:

Administrative client macros

MIGRATE STGPOOL (Migrate storage pool to next storage pool)

Use this command to migrate files from one storage pool to the next storage pool in the storage hierarchy.

This command can only be used with primary storage pools. The storage pool data format cannot be NETAPPDUMP, CELERRADUMP, or NDMPDUMP. Data cannot be migrated into or out of storage pools that are defined with a CENTERA device class.

Only one migration or reclamation process for a given storage pool is allowed at any given time. If a migration or reclamation process is already running for the storage pool, you cannot start another migration process for the storage pool.

You should only use this command if you are not going to use automatic migration for the storage pool. To prevent automatic migration from running, set the HIGHMIG attribute of the storage pool definition to 100.

If you use this command to start a migration process, but the storage pool does not have a next storage pool identified in the hierarchy, a reclamation process is triggered for the source storage pool. To prevent the reclamation process, define the next storage pool in the hierarchy. Then, start the migration process.

The MIGRATE STGPOOL command honors the values of the following parameters on the DEFINE STGPOOL and UPDATE STGPOOL commands:

- MIGPROCESS
- MIGDELAY
- MIGCONTINUE
- NEXTPOOL
- LOWMIG

Tip: You can override the value of the LOWMIG parameter on DEFINE STGPOOL and UPDATE STGPOOL by specifying a value for the LOWMIG parameter on the MIGRATE STGPOOL command.

The MIGRATE STGPOOL command ignores the value of the HIGHMIG parameter of the storage pool definition. Migration occurs regardless of the value of the HIGHMIG parameter.

This command creates one or more migration processes that can be canceled with the CANCEL PROCESS command. The number of processes is limited by the MIGPROCESS attribute of the storage pool definition. To display information about background

processes, use the QUERY PROCESS command.

Remember: Migrating data from a primary storage pool that is set up for data deduplication to another primary storage pool that is also set up for data deduplication removes duplicate data.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for both the storage pool from which the files are to be migrated and the next storage pool to which files are to be migrated.

Syntax

```
>>-MIGrate STGpool--pool_name--+-----+----->
                                     '-LOWmig----number-'
                                     .-REClaim---No-----
>--+-----+-----+----->
   '-DUration---minutes-'   '-REClaim---+No---+'
                               '-Yes-'

   .-Wait-----No-----
>--+-----+-----><
   '-Wait-----+No---+'
                               '-Yes-'
```

Parameters

pool_name (Required)

Specifies the primary storage pool from which files are to be migrated.

DUration

Specifies the maximum number of minutes the migration runs before being automatically canceled. When the specified number of minutes elapses, the server will automatically cancel all migration processes for this storage pool. As soon as the processes recognize the automatic cancellation, they end. As a result, the migration might run longer than the value you specified for this parameter. You can specify a number from 1 to 9999. This parameter is optional. If not specified, the server will stop only after the low migration threshold is reached.

LOWmig

For random-access and sequential-access disk storage pools, specifies that migration should stop when the amount of data in the pool is at or below this percentage of the pool's estimated capacity. This parameter is optional.

The calculation for sequential-access disk storage pools includes the capacity of all the scratch volumes that are specified for the pool. Because migration is by node or filesystem, depending upon collocation, the occupancy of the storage pool can fall below the value that you specified for this parameter. To empty the storage pool, set LOWMIG=0. For other types of sequential-access storage pools, the server stops migration when the ratio of volumes containing data to the total number of volumes in the storage pool is at or below this percentage. The total number of volumes includes the maximum number of scratch volumes. You can specify a number from 0 to 99 for this optional parameter. The default value is the LOWMIG attribute of the storage pool definition.

REClaim

Specifies whether reclamation is attempted for the storage pool before completing the migration. This parameter can only be specified for a sequential-access storage pool. This parameter is optional. The default is No. Possible values are:

No

Specifies that the server will not attempt a reclamation before starting the migration.

Yes

Specifies that the server will attempt reclamation before starting the migration. Any volumes in the storage pool that meet the reclamation threshold as specified by the RECLAIM attribute of the storage pool definition will be reclaimed before completing the migration. If no volumes meet the reclamation threshold or if, after reclamation, the LOWMIG threshold has not been reached, the server will begin the migration. Before reclaiming space for storage pools defined with RECLAMATIONTYPE=SNAPLOCK, the server deletes all empty WORM FILE volumes during reclamation processing that have exceeded their reclaim period.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. This default is No. Possible values are:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files may have already been migrated before the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where the messages are logged.

Note: You cannot specify WAIT=YES from the server console.

Example: Migrate a storage pool to the next storage pool

Migrate data from the storage pool named BACKUPPOOL to the next storage pool. Specify that the server should end the migration as soon as possible after 90 minutes.

```
migrate stgpool backuppool duration=90
```

Related commands

Table 1. Commands related to MIGRATE STGPOOL

Command	Description
CANCEL PROCESS	Cancel a background server process.
QUERY PROCESS	Displays information about background process.
QUERY STGPOOL	Displays information about storage pools.
RECLAIM STGPOOL	Performs reclamation for the storage pool.

Related information:

[Migrating files in a storage pool hierarchy](#)

MOVE commands

Use the MOVE commands to either transfer backup or archive data between storage pools, or to move disaster recovery media on and off site.

- MOVE CONTAINER (Move a container)
- MOVE DATA (Move files on a storage pool volume)
- MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)
- MOVE GRPMEMBER (Move a server group member)
- MOVE MEDIA (Move sequential-access storage pool media)
- MOVE NODEDATA (Move data by node in a sequential access storage pool)

AIX

Linux

Windows

MOVE CONTAINER (Move a container)

Use this command to move the contents of a storage pool container to another container if a storage pool directory is removed or if a container is damaged. You can also use the command to consolidate data and reclaim space. You can issue this command for directory containers and cloud containers.

If the data in a storage pool is fragmented, the command consolidates the data:

- For a directory-container storage pool, the command potentially reduces the number of containers.
- For a cloud-container storage pool, the command consolidates the data into a smaller container.

In addition, for directory-container storage pools, you can use this command to move the contents of a storage pool container under these conditions:

- When you upgrade hardware
- If I/O errors occur on a disk

Privilege class

To issue this command, you must have restricted storage privilege.

Syntax

```

                                .-DEFrag----Yes-----.
>>-MOVE CONTainer--container_name--+-----+----->
                                '-DEFrag----+-Yes+-'
                                    '-No--'

>--+-----+----->
    '-STGPOOLDIRectory----directory_name-'

    .-Wait----Yes-----.
>--+-----+-----><
    '-Wait----+-Yes+-'
        '-No--'
```

Parameters

container_name (Required)

Specifies the name of the container to move. You must specify the full path name of the container.

DEFrag

Specifies whether the contents of the container are consolidated into existing containers during a MOVE CONTAINER operation. This parameter is optional.

The following values are possible:

Yes

This is the default value. The container contents are moved in the following way:

- For a container in a directory-container storage pool, the contents are moved into one or more existing containers. If the existing containers have insufficient space, a container is created and any remaining data is allocated to the new container.
- For a container in a cloud-container storage pool, the contents are moved into a single new cloud container.

No

The contents are moved into a newly created container.

Restriction: If you are issuing the MOVE CONTAINER command for a cloud container, you cannot specify DEFrag=NO.

In some cases, especially if you encrypt data, you might have to create additional containers and allocate the data to the new containers to ensure sufficient space. For instructions, see technote 7050411.

STGPOOLDIRectory

Specifies the name of the storage pool directory to which the container is moved. This parameter is optional.

If you specify a storage pool directory, it must be in the same storage pool as the original container. The storage pool directory is used for the new container. If you don't specify a storage pool directory, the IBM Spectrum Protect™ server selects a storage pool directory from the same storage pool.

Restriction: If you are issuing the MOVE CONTAINER command for a cloud container, do not specify the STGPOOLDIRectory parameter.

Wait

Specifies whether to wait for the IBM Spectrum Protect server to process this command in the foreground. This parameter is optional. You can specify one of the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages that are related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged. This is the default.

Yes

The server processes this command in the foreground. The operation must complete processing before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify the WAIT=YES parameter from the server console.

Example: Move a container in a directory-container storage pool

AIX | **Linux** Move a container, 0000000000000001.dcf, from the /data1/storage/dir1 storage pool directory to the /data/storage/dir2 storage pool directory.

```
move container /data1/storage/dir1/00/0000000000000001.dcf
stgpooldir=/data/storage/dir2
```

Windows Move a container, 0000000000000001.dcf, from the e:\data1\storage\dir1 storage pool directory to the e:\data\storage\dir2 storage pool directory.

```
move container e:\data1\storage\dir1\00\0000000000000001.dcf
stgpooldir=e:\data\storage\dir2
```

Table 1. Commands related to MOVE CONTAINER

Command	Description
AUDIT CONTAINER commands	Audit directory-container or cloud-container storage pools.
QUERY CONTAINER	Displays information about a container.

MOVE DATA (Move files on a storage pool volume)

Use this command to move files from one storage pool volume to other storage pool volumes.

Restriction: You cannot use this command for volumes that are assigned to copy-container storage pools.

You can move files from a primary storage pool volume only to volumes in the same or a different primary storage pool. You can move files from a copy storage pool volume only to volumes in the same copy storage pool. You can move files from an active-data pool volume only to volumes in the same active-data pool.

In addition to moving data from volumes in storage pools that have NATIVE or NONBLOCK data formats, you can use this command to move data from volumes in storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The target storage pool must have the same data format as the source storage pool. If you are moving data out of a storage pool for the purpose of upgrading to new tape technology, the target primary storage pool must be associated with a library that has the new device for the tape drives. IBM Spectrum Protect™ supports backend data movement for NDMP images.

You cannot move data into or out of a storage pool that is defined with a CENTERA device class.

If you are moving files to volumes in the same storage pool, sufficient space must be available on the volumes. Otherwise, the operation fails.

When you move files from a sequential access volume, multiple sequential access volume mounts are required to move files that span volumes.

When you move files from a random access volume, the server erases any cached copies of files on the volume.

After a move data operation completes, a volume might not be empty if one or more files cannot be relocated to another volume because of input/output errors on the device or because errors were found in the file. If needed, you can delete the volume using the option to discard any data. The files with I/O or other errors are then deleted.

You can use this command to move files from an offsite volume in a copy storage pool or active-data pool. Because the offsite volume cannot be mounted, the server obtains the files that are on the offsite volume from either a primary storage pool or another copy storage pool. These files are then written to the destination volumes in the original copy storage pool or active-data pool.

During the data movement process, active-data pools cannot be used to obtain data.

If you run the MOVE DATA command on an offsite volume that contains collocated data, it might be necessary to issue the MOVE DATA command multiple times to move all of the data out of the volume. For example, if you are using filespace collocation groups with an offsite volume that contains filespace in a collocation group and filespace that are not in the group, you must issue two MOVE DATA commands. Each MOVE DATA command moves the data for a single collocated or non-collocated group of files.

Do not use the MOVE DATA command if a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The MOVE DATA command might cause the restore to be incomplete. If you issue the MOVE DATA command during a restore operation and you receive an error message indicating that one or more files are locked and cannot be moved, you must reissue the MOVE DATA command after the restore operation completes in order to move any remaining files.

Remember:

Issuing this command removes duplicate data when:

- Moving data from a primary storage pool that is set up for data deduplication to another primary storage pool that is also set up for data deduplication.
- Moving data within a copy storage pool that is set up for data deduplication.
- Moving data within an active-data pool that is set up for data deduplication.

A volume in a deduplicated storage pool might contain files that are logically deleted but are still linked by files on other volumes. If you use the MOVE DATA command to move the contents of a deduplicated storage pool volume to a non-deduplicated storage pool, the logically deleted files are not written to the new volume since they do not exist logically. The deleted files are kept on the original volumes for other files to reference. The MOVE DATA process ends successfully but none of the deleted files are moved to the new target volume and the source volume is not deleted. You can issue the QUERY CONTENT command with the FOLLOWLINKS=YES or FOLLOWLINKS=JUSTLINKS parameter to verify whether the volume contains files that are linked by files on other volumes.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume belongs and also for the new storage pool, if one is specified.

Syntax

```
>>-MOVE Data--volume_name--+-----+----->
                               '-STGpool---pool_name-'

.-SHREDTONOshred---No-----
>--+-----+----->
   '-SHREDTONOshred---+No--+-'
                               '-Yes-'

                               (1) (2)
.-RECONStruct---No or Yes-----
>--+-----+----->
   '-RECONStruct---+No--+-'
                               '-Yes-'

.-Wait---No-----
>--+-----+-----><
   '-Wait---+No--+-'
                               '-Yes-'
```

Notes:

1. The default is NO if either the source or target storage pool is random access. The default is YES if both the source and target storage pools are sequential access.
2. This parameter is not available or is ignored if the data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP data.

Parameters

volume_name (Required)
Specifies the storage pool volume from which to move files.
STGpool

Specifies the primary storage pool to which you want to move files (the target storage pool). This parameter is optional and applies only to moving data from primary storage pool volumes. If you do not specify a value for this parameter, files are moved to other volumes within the same storage pool.

SHREDTONOshred

Specifies whether data is moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the server will not allow data to be moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. If the source storage pool enforces shredding and the target storage pool does not, the operation fails.

Yes

Specifies that the server allows data to be moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. The source data is shredded when the operation is complete. The target data will not be shredded when it is deleted.

RECONStruct

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that has accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when moving the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

Possible values are:

No

Specifies that reconstruction of file aggregates is not completed during data movement.

Yes

Specifies that reconstruction of file aggregates is completed during data movement. You can only specify this option when both the source and the target storage pools are sequential-access.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If a MOVE DATA background process is canceled, some files may have already moved before the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Move files on a storage pool volume

Move files from storage pool volume STGVOL.1 to any available volumes assigned to the 8MMPOOL storage pool.

```
move data stgvol.1 stgpool=8mmpool
```

Related commands

Table 1. Commands related to MOVE DATA

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
SHRED DATA	Manually starts the process of shredding deleted data.

MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)

Use this command to track volumes that are to be moved offsite and to identify the expired or empty volumes that are to be moved onsite. You can track database backup volumes, and volumes in copy storage pools, container-copy storage pools, and active-data storage pools.

The processing of volumes by this command depends on what the volumes are used for:

Backups of the server database

To control whether the command processes database backup volumes, use the SOURCE parameter on this command. The command can process volumes that are used for full plus incremental or snapshot database backups. You cannot specify virtual volumes (backup objects that are stored on another server). You can change volumes through each state, or you can use the TOSTATE parameter and skip states to simplify the movements.

Copy storage pools

The MOVE DRMEDIA command always processes copy storage-pool volumes.

Container-copy storage pools

By default, volumes in container-copy storage pools are not eligible for processing by the MOVE DRMEDIA command. To process container-copy storage pool volumes, you must issue the SET DRMCOPYCONTAINERSTGPOOL command first, or specify the COPYCONTAINERSTGPOOL parameter on the MOVE DRMEDIA command.

Active-data storage pools

By default, volumes in active-data storage pools are not eligible for processing by the MOVE DRMEDIA command. To process active-data pool volumes, you must issue the SET DRMACTIVEDATASTGPOOL command first, or specify the ACTIVEASTGPOOL parameter on the MOVE DRMEDIA command.

You can use the QUERY ACTLOG command to see whether the MOVE DRMEDIA command was successful. You can also view this information from the server console.

Restriction: Do not run the MOVE DRMEDIA and BACKUP STGPOOL commands concurrently. Ensure that the storage pool backup processes are complete before you issue the MOVE DRMEDIA command.

Privilege class

To issue this command, you must have one of the following privilege classes:

- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO: operator, unrestricted storage, or system privilege.
- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES (the default): system privilege.

Syntax

```

>>-MOVE DRMedia--volume_name----->
>--+-----+----->
  '-WHEREState---+-Mountable-----+'
      +-NOTMountable-----+
      +-COUrier-----+
      +-VAULTRetrieve---+
      '-COURIERRetrieve-'
>--+-----+----->
  '-BEGINdate---date-' '-ENDDate---date-'
>--+-----+----->
  '-BEGINtime---time-' '-ENDTime---time-'
>--+-----+----->
  '-COPYCONtainerstgpool---pool_name-'
>--+-----+----->
  '-COPYstgpool---pool_name-'
>--+-----+----->
  '-ACTIVEDatastgpool---pool_name-'
  .-Source---DBBackup-----
>--+-----+----->
  '-Source---DBBackup---+'
      +-DBSnapshot-+
      '-DBNOne-----'
  .-REMove---Bulk-----
>--+-----+----->
  '-REMove---+-No-----+'
      +-Yes-----+
      +-Bulk-----+
      '-Untileefull-'
>--+-----+----->
  '-TOSTate---+-NOTMountable-----+'
      +-COUrier-----+
      +-VAult-----+
      +-COURIERRetrieve-+
      '-ONSITERetrieve--'
>--+-----+----->
  '-WHERELOcation---location-'
>--+-----+----->
  '-TOLOcation---location-' '-Cmd---"command"-
  .-APPend---No-----
>--+-----+----->
  '-CMDFilename---file_name-' '-APPend---+-No--+-'
      '-Yes-'
  .-Wait---No-----
>--+-----+----->>
  '-Wait---+-No--+-' '-CAP---x,y,z-'
      '-Yes-'

```

Parameters

volume_name (Required)

Specifies the name of the volume to be processed. You can use wildcard characters. If you use wildcard characters to specify this name, you must also specify the WHERESTATE parameter. The server looks for matching names among the following eligible volumes:

- Database backup volumes, as specified by the SOURCE parameter of this command.
- Copy storage pool volumes from the storage pools named in the COPYSTGPOOL parameter. If you do not use the COPYSTGPOOL parameter, the server processes volumes from copy storage pools that were previously specified in

the SET DRMCOPYSTGPOOL command.

- Container-copy storage pool volumes from the storage pools named in the COPYCONTAINERSTGPOOL parameter. If you do not use the COPYCONTAINERSTGPOOL parameter, the server processes volumes from container-copy storage pools that were previously specified in the SET DRMCOPYCONTAINERSTGPOOL command.
- Active-data storage pool volumes from the storage pools named in the ACTIVEDATASTGPOOL parameter. If you do not use the ACTIVEDATASTGPOOL parameter, the server processes volumes from active-data storage pools that were previously specified in the SET DRMACTIVEDATASTGPOOL command.

Other parameters can also limit the results of the command.

WHEREState

Specifies the state of volumes to be processed. This parameter is required if the TOSTATE parameter is not specified or if you use a wildcard character in the volume name. For more information, see Table 2 and Table 3. Specify one of the following values:

MOutable

These volumes contain valid data and are available for onsite processing. The values change to NOTMOUNTABLE if the TOSTATE parameter is not specified.

Depending on the outcome of the REMOVE parameter, the server might eject volumes in an automated library before you change the destination state.

For external libraries, the server sends requests to the external library manager to eject the volumes. It depends on the external library manager whether the volumes are ejected from the library.

NOTMOutable

These volumes are onsite, contain valid data, and are not available for onsite processing. The values change to COURIER if the TOSTATE parameter is not specified.

COURier

These volumes are with the courier and being moved offsite. The values change only to VAULT.

VAULTRetrieve

These volumes are at the offsite vault and do not contain valid data. The values change to COURIERRETRIEVE if the TOSTATE parameter is not specified.

COURIERRetrieve

These volumes are with the courier and being moved onsite. The values change only to ONSITERETRIEVE. The server deletes the volume records of the database backup and scratch copy storage pool volumes from the database.

BEGINDate

Specifies the beginning date that is used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command changes the volume to its current state on or after the specified date. The default is the earliest date for which volume information exists.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/1998
TODAY	The current date.	TODAY
TODAY-days or -days	The current date minus days specified.	TODAY-7 or -7 To identify volumes that were changed to their current state a week ago, you can specify TODAY-7 or -7.
EOLM (end of last month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (beginning of this month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies the ending date that is used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command changes the volume to its current state on or before the specified date. The default is the current date.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/1998
TODAY	The current date.	TODAY To identify volumes that were changed to their current state today, specify TODAY.
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY-1 or -1 To identify volumes that were changed to their current state a week ago, you can specify TODAY-1 or -1.
EOLM (end of last month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (beginning of this month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time that is used to select volumes for processing. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command changes the volume to its current state on or after the specified time and date. The default is midnight (00:00:00) on the date that is specified with the BEGINDATE parameter.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date.	12:33:28
NOW	The current time on the specified begin date.	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date.	NOW+03:00 or +03:00
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date.	NOW-03:30 or -03:30 If you issue the MOVE DRMEDIA command at 9:00 with BEGINTIME=NOW-03:30 or BEGINTIME=-03:30, the server identifies the volumes that were changed to their current state at 5:30 on the begin date that you specify.

ENDTime

Specifies the ending time that is used to select volumes for processing. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command changes the volume to its current state on or after the specified time and date. The default is 23:59:59.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date.	12:33:28
NOW	The current time on the specified end date.	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date.	NOW+03:00 or +03:00 If you issue the MOVE DRMEDIA command at 9:00 with ENDTIME=NOW+03:30 or ENDTIME=+03:30, the server identifies the volumes that were changed to their current state at 12:30 on the end date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date.	NOW-03:30 or -03:30

COPYCONTAINERSTGPPOOL

Specifies the name of the container-copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. If you use wildcard characters to specify this name, you must also specify the WHERESTATE parameter.

The container-copy storage pools that are specified with this parameter override storage pools that are specified with the SET DRMCOPYCONTAINERSTGPPOOL command. If this parameter is not specified, the server selects the storage pools as follows:

- If the SET DRMCOPYCONTAINERSTGPPOOL command was previously issued with valid container-copy storage pool names, the server processes only those storage pools.
- If the SET DRMCOPYCONTAINERSTGPPOOL command was not issued, or if all of the container-copy storage pools were removed by using the SET DRMCOPYCONTAINERSTGPPOOL command, the server processes all container-copy storage pool volumes based on the setting of the WHERESTATE parameter. If the parameter is set to a value of NOTMOUNTABLE, COURIER, VAULTRETRIEVE, or COURIERRETRIEVE, the volumes are processed. If the value is MOUNTABLE, the volumes are not processed.

COPYSTGPPOOL

Specifies the name of the copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. If you use wildcard characters to specify this name, you must also specify the WHERESTATE parameter.

The copy storage pools that are specified with this parameter override copy storage pools that are specified with the SET DRMCOPYSTGPPOOL command. If this parameter is not specified, the server selects the storage pools as follows:

- If the SET DRMCOPYSTGPPOOL command was previously issued with valid copy storage pool names, the server processes only those storage pools.
- If the SET DRMCOPYSTGPPOOL command was not issued, or if all of the copy storage pools are removed by using the SET DRMCOPYSTGPPOOL command, the server processes all copy storage pool volumes in the specified state. The states available are MOUNTABLE, NOTMOUNTABLE, COURIER, VAULTRETRIEVE, or COURIERRETRIEVE.

ACTIVEDATASTGPPOOL

Specifies the name of the active-data pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. If you use wildcard characters to specify this name, you must also specify the WHERESTATE parameter.

The active-data pools that are specified with this parameter override active-data pools that are specified with the SET DRMACTIVEDATASTGPPOOL command. If this parameter is not specified, the server selects the storage pools in the following way:

- If the SET DRMACTIVEDATASTGPPOOL command was previously issued with valid active-data pool names, the server processes only those storage pools.
- If the SET DRMACTIVEDATASTGPPOOL command was not issued, or all of the active-data pools are removed by using the SET DRMACTIVEDATASTGPPOOL command, the server processes all active-data pool volumes in the specified state. The states available are NOTMOUNTABLE, COURIER, VAULTRETRIEVE, or COURIERRETRIEVE. Volumes in the MOUNTABLE state are not processed.

Source

Specifies whether to include database backup volumes for processing. This parameter is optional. The default is DBBACKUP. Specify one of the following values:

DBBackup

Specifies that the server includes full and incremental database backup volumes for processing.

DBSnapshot

Specifies that the server includes database snapshot backup volumes for processing.

DBNone

Specifies that the server does not include any database backup volumes for processing.

REMove

Specifies that the server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values are YES, NO, BULK, and UNTILEEFULL. The default is BULK. The response of the server to each value and the default value depends on the type of library.

Restriction: You can use the REMOVE=UNTILEEFULL option only with the library type SCSI.

SCSI libraries

The response of the server to the command depends on whether the library has entry/exit ports, and if so, whether a port is available for use. See the following table.

Table 1. Server response for SCSI libraries

Library characteristic	Server response when you specify REMOVE=YES	Server response when you specify REMOVE=BULK	Server response when you specify REMOVE=NO	Server response when you specify REMOVE=UNTILEEFULL
Library has no entry/exit ports	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
Library has entry/exit ports and an entry/exit port is available	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.

Library characteristic	Server response when you specify REMOVE=YES	Server response when you specify REMOVE=BULK	Server response when you specify REMOVE=NO	Server response when you specify REMOVE=UNTILE EFULL
Library has entry/exit ports, but no ports are available	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server waits for a port to be made available.	The server specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The command fails and any remaining eligible volumes are not processed. Make the port available and issue the command again.

349X libraries

REMOVE=YES

The 3494 Library Manager ejects the cartridge to the convenience I/O station.

REMOVE=BULK

The 3494 Library Manager ejects the cartridge to the high-capacity output facility.

REMOVE=NO

The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the INSERT category for use by other applications.

ACSLs libraries

REMOVE=YES or REMOVE=BULK

The server ejects the cartridge to the convenience I/O station.

The server then deletes the volume entry from the server library inventory.

When you move volumes from the MOUNTABLE state with REMOVE=YES specified, the MOVE MEDIA command uses more than one slot in the CAP for a StorageTek library with ACSLS.

REMOVE=NO

The server does not eject the cartridge.

The server deletes the volume entry from the server library inventory and leaves the volume in the library.

External libraries

You can specify REMOVE=YES, REMOVE=BULK, or REMOVE=NO. For any value, the server requests the external library manager to eject the volume from the library.

It depends on the external library manager whether the volume is ejected from the library. Refer to the external library documentation for information about the procedures to follow when you use the MOVE DRMEDIA command to track volumes.

TOSTate

Specifies the destination state of the volumes that are processed. This parameter is required if the WHERESTATE parameter is not specified. If you specify TOSTATE parameter but not WHERESTATE parameter, you must specify the volume name. Wildcard characters are not allowed. See Table 2 and Table 3.

Specify one of the following values:

NOTMOUNTable

Specifies that volumes are to change to the NOTMOUNTABLE state. This value is valid only if the volumes are in the MOUNTABLE state.

If volumes are in an automated library, the server might eject the volumes from the library before you change them to the NOTMOUNTABLE state, depending on the behavior of the REMOVE parameter.

For external libraries, the server sends requests to the external library manager to eject the volumes. Whether the volumes are ejected from the library depends on the external library manager. Refer to the external library documentation for information about the procedures to follow when you use the MOVE DRMEDIA command to track the volumes.

COURier

Specifies that volumes are to change to the COURIER state. This value is valid only if the volumes are in the MOUNTABLE or NOTMOUNTABLE state.

Depending on the behavior of the REMOVE parameter and whether volumes are in an automated library, the server might eject the volumes from the library before you change them to the COURIER state.

For external libraries, the server sends requests to the external library manager to eject the volumes. Whether the volumes are ejected from the library depends on the external library manager. Refer to the external library documentation for information about the procedures to follow when you use the MOVE DRMEDIA command to track the volumes.

VAult

Specifies that volumes are to change to the VAULT state. This value is valid only if the volumes are in the MOUNTABLE, NOTMOUNTABLE, or COURIER state.

Depending on the behavior of the REMOVE parameter and whether volumes are in an automated library, the server might eject the volumes from the library before you change them to the VAULT state.

For external libraries, the server sends requests to the external library manager to eject the volumes. Whether the volumes are ejected from the library depends on the external library manager. Refer to the external library documentation for information about the procedures to follow when you use the MOVE DRMEDIA command to track the volumes.

COURIERRetrieve

Specifies that volumes are to change to the COURIERRETRIEVE state. This value is valid only if the volumes are in the VAULTRETRIEVE state.

ONSITERetrieve

Specifies that volumes are to change to the ONSITERETRIEVE state. This value is valid only if the volumes are in the VAULTRETRIEVE or COURIERRETRIEVE state. For database backup and scratch copy storage pool volumes that are changing to the ONSITERETRIEVE state, the server deletes the volume records from the database.

WHERELocation

Specifies the current location of the volumes. This parameter is optional. The maximum length of the location is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

TOLOcation

Specifies the destination location of the volumes. This parameter is optional. The maximum length of the location that is specified is 255 characters. Enclose the text in quotation marks if it contains any blank characters. If you do not specify the destination location, the location that is defined by the SET DRMNOTMOUNTABLE command is used.

CMd

Specifies a command to be issued for each volume that is processed by the MOVE DRMEDIA command. DRM writes the commands to a file that is specified by the CMDFILENAME parameter. After the MOVE DRMEDIA operation is completed, the commands in the file can be issued. The command can contain up to 255 characters. If the command contains more than 240 characters, it is split into multiple lines, and continuation characters (+) are added. You might need to alter the continuation character based on the operating system. This parameter is optional.

command

The command string that is enclosed in quotation marks. The string must not include embedded quotation marks. For example, the following CMD parameter is valid:

```
cmd="checkin libvol lib8mm &vol status=scratch"
```

The following example is not a valid way to specify the CMD parameter:

```
cmd=""checkin libvol lib8mm" &vol status=scratch""
```

The command can include substitution variables. The variables are not case-sensitive, and must not contain blank spaces after the ampersand (&). You can specify the following values:

&VOL

A volume name.

&LOC

A volume location.

&VOLDSN

The file name to be written into the sequential access media labels. For example, if the applicable device class sets BKP as the tape volume prefix, a copy storage pool tape volume file name might be BKP.BFS and a database backup tape volume file name might be BKP.DBB.

&NL

The new line character. When you use the new line character, the command is split at the &NL variable. If required, you must specify the appropriate continuation character before the &NL character. If the &NL character is not specified and the command line is greater than 240 characters, the line is split into multiple lines and continuation characters (+) are added.

AIX Linux CMDFilename

AIX Linux Specifies the fully qualified name of the file that contains the commands that are specified by CMD parameter. This parameter is optional.

If you do not specify a file name or if you specify a null string (""), DRM uses the file name that is specified by the SET DRMCMDFILENAME command. If you do not specify a file name with the SET DRMCMDFILENAME command, DRM generates a file name by appending `exec.cmds` to the directory path name of the current working directory of the server.

If the operation fails after the command file is created, the file is not deleted.

Windows CMDFilename

Windows Specifies the fully qualified name of the file that contains the commands that are specified by CMD parameter. This parameter is optional.

The maximum length of the file name is 259 characters. If you do not specify a file name or if you specify a null string (""), DRM uses the file name that is specified by the SET DRMCMDFILENAME command. If you do not specify a file name with the SET DRMCMDFILENAME command, DRM generates a file name by appending `exec.cmd` to the directory that represents this instance of the server (typically the directory from which the server was installed). The DRM allocates the file name that is specified or generated. If the file name exists, DRM tries to use it; any existing data is overwritten. If this happens and the executable commands in the file have not been run, issue QUERY DRMEDIA command to rebuild the executable commands for the desired date and volume transition.

If the MOVE DRMEDIA command fails and none of the command string that is specified with the CMD parameter is written for the volume that successfully moved, the allocated file name is deleted.

APPend

Specifies whether to overwrite any existing contents of the command file or append the commands to the file. This parameter is optional. The default is NO. Specify one of the following values:

No

DRM overwrites the contents of the file.

Yes

DRM appends the commands to the file.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. Specify one of the following values:

No

Specifies that the server processes this command in the background.

Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To see whether the operation was successful, issue the QUERY ACTLOG command.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before you continue with other tasks. The server then displays the output messages to the administrative client.

Restriction: You cannot specify WAIT=YES from the server console.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (x,y,z), issue the QUERY CAP command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

x

The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.

y

The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.

z

The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Rules for destination states and destination locations

The following table shows how DRM determines the destination state and location of a volume.

Destination state

- The value of the TOSTATE parameter that was specified
- The next state of the WHERESTATE parameter that was specified, if the TOSTATE parameter was not specified

Destination location

- The value of the TOLOCATION parameter that was specified
- The location of the TOSTATE parameter that was specified, if the TOLOCATION parameter was not specified
- The location of the next state of the WHERESTATE parameter that was specified, if the TOLOCATION and TOSTATE parameters are not specified

Table 2. Volume destination and location

Parameters specified	Destination state	Destination location
WHERESTATE	The next state of the WHERESTATE	Location of the next state
WHERESTATE, TOSTATE	TOSTATE	Location of the TOSTATE
WHERESTATE, TOLOCATION	The next state of the WHERESTATE	TOLOCATON
WHERESTATE, TOSTATE, TOLOCATION	TOSTATE	TOLOCATION
TOSTATE	TOSTATE	Location of the TOSTATE
TOSTATE, WHERELOCATION	TOSTATE	Location of the TOSTATE
TOSTATE, WHERELOCATION, TOLOCATION	TOSTATE	TOLOCATION

Rules for state transitions

The following tables show the state transitions that volumes are eligible for, based on their current state.

Table 3. State transitions for volumes

The current state of the volume	Destination state		
	MOUNTABLE	NOTMOUNTABLE	COURIER
MOUNTABLE	N	Y	Y
NOTMOUNTABLE	N	N	Y
COURIER	N	N	N
VAULT	N	N	N
VAULTRETRIEVE	N	N	N

The current state of the volume	Destination state		
	MOUNTABLE	NOTMOUNTABLE	COURIER
COURIERRETRIEVE	N	N	N
ONSITERETRIEVE	N	N	N

Table 4. State transitions for volumes

The current state of the volume	Destination state	
	VAULT	VAULTRETRIEVE
MOUNTABLE	Y	N
NOTMOUNTABLE	Y	N
COURIER	Y	N
VAULT	N	N
VAULTRETRIEVE	N	N
COURIERRETRIEVE	N	N
ONSITERETRIEVE	N	N

Table 5. State transitions for volumes

The current state of the volume	Destination state	
	COURIERRETRIEVE	ONSITERETRIEVE
MOUNTABLE	N	N
NOTMOUNTABLE	N	N
COURIER	N	N
VAULT	N	N
VAULTRETRIEVE	Y	Y
COURIERRETRIEVE	N	Y
ONSITERETRIEVE	N	N

Example: Move disaster recovery media from the NOTMOUNTABLE state

Move disaster recovery media that is in the NOTMOUNTABLE state to the COURIER state, and then query the results.

```
move drmedia * wherestate=notmountable
tostate=courier
```

```
query actlog search="MOVE DRMEDIA"
```

```
08/11/1999 11:12:24 ANR0984I Process 10 for MOVE DRMEDIA started
in the BACKGROUND at 11:12:24.
08/11/1999 11:12:24 ANR0610I MOVE DRMEDIA started by HSIAO as
process 10.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume TAPE0P was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume TAPE1P was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume DBTP02 was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume DBTP01 was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6682I MOVE DRMEDIA command ended: 4 volumes
processed.
08/11/1999 11:12:25 ANR0611I MOVE DRMEDIA started by HSIAO as
process 10 has ended.
08/11/1999 11:12:25 ANR0985I Process 10 for MOVE DRMEDIA running in
the BACKGROUND processed 4 items with a
completion state of SUCCESS at 11:12:25.
```

Example: Move disaster recovery media from the MOUNTABLE state

Move disaster recovery media from the MOUNTABLE state to the COURIER state. If the media is in an automated library, MOVE DRMEDIA ejects the media before you change the state.

```
move drmedia * wherestate=mountable tostate=courier wait=yes
```

```
ANR0984I Process 12 for MOVE DRMEDIA started
  in the FOREGROUND at 09:57:17.
ANR0609I MOVE DRMEDIA started as process 12.
ANR0610I MOVE DRMEDIA started by HSIAO as
  process 12.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume TAPE01 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume TAPE01 in library LIB8MM completed
  successful.
ANR6683I MOVE DRMEDIA: Volume TAPE01 was moved
  from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume TAPE02 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume TAPE02 in library LIB8MM completed
  successful.
ANR6683I MOVE DRMEDIA: Volume TAPE02 was moved
  from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume DBTP05 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume DBTP05 in library LIB8MM completed
  successful.
ANR6683I MOVE DRMEDIA: Volume DBTP05 was moved
  from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume DBTP04 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
  volume DBTP04 in library LIB8MM completed
  successful.
ANR6683I MOVE DRMEDIA: Volume DBTP04 was moved
  from MOUNTABLE state to COURIER.
ANR6682I MOVE DRMEDIA command ended: 4 volumes
  processed.
ANR0611I MOVE DRMEDIA started by HSIAO as
  process 12 has ended.
ANR0985I Process 12 for MOVE DRMEDIA running
  in the FOREGROUND processed 4 items with a
  completion state of SUCCESS at 10:12:25.
```

Example: Move disaster recovery media from the VAULTRETRIEVE state

Move disaster recovery media that is in the VAULTRETRIEVE state to the ONSITERETRIEVE state. Generate a CHECKIN LIBVOLUME command for each volume that is successfully processed and store the commands in a file:

AIX | **Linux**

```
move drmedia * wherestate=vaultretrieve tostate=onsiteretrieve
cmdfilename=/drm/move/exec.cmds
cmd="checkin libvol lib8mm &vol status=scratch"
```

Windows

```
move drmedia * wherestate=vaultretrieve tostate=onsiteretrieve
cmdfilename=c:\drm\move\exec.cmd
cmd="checkin libvol lib8mm &vol status=scratch"
```

Query the results:

```
query actlog search="MOVE DRMEDIA"

08/13/1999 09:12:24 ANR0984I Process 15 for MOVE DRMEDIA started in
                    the BACKGROUND at 09:12:24.
08/13/1999 09:12:24 ANR0610I MOVE DRMEDIA started by HSIAO as
                    process 15.
```

```

08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume CSTEP01 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume CSTEP02 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume DBTP10 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume DBTP11 was deleted.
08/13/1999 09:12:27 ANR6682I MOVE DRMEDIA command ended: 4 volumes
                        processed.
08/13/1999 09:12:42 ANR0611I MOVE DRMEDIA started by HSIAO as process
                        15 has ended.
08/13/1997 09:12:42 ANR0985I Process 15 for MOVE DRMEDIA running in
                        the BACKGROUND processed 4 items with a
                        completion state of SUCCESS at 09:12:42.

```

The volume check-in commands were also created in the file that was specified with the CMDFILENAME parameter:

- **AIX** | **Linux** /drm/move/exec.cmds
- **Windows** c:\drm\move\exec.cmd

The file contains these lines:

```

checkin libvol lib8mm CSTEP01 status=scratch
checkin libvol lib8mm CSTEP02 status=scratch
checkin libvol lib8mm DBTP10 status=scratch
checkin libvol lib8mm DBTP11 status=scratch

```

Tip: To process the CHECKIN LIBVOLUME commands, issue the MACRO command with the file name as the macro name.

Related commands

Table 6. Commands related to MOVE DRMEDIA

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CANCEL PROCESS	Cancels a background server process.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
PREPARE	Creates a recovery plan file.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY PROCESS	Displays information about background processes.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
AIX Linux Windows SET DRMCOPYCONTAINERSTGPOOL	AIX Linux Windows Specifies the container-copy storage pools that are used in DRM commands.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMCOURIERNAME	Specifies the name of the courier for the disaster recovery media.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.
SET DRMVAULTNAME	Specifies the name of the vault where DRM media is stored.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.

Command	Description
SET DRMNOTMOUNTABLENAME	Specifies the location name of the DRM media to be sent offsite.

MOVE GRPMEMBER (Move a server group member)

Use this command to move a member from one server group to another server group. The command fails if the member you are moving has the same name as a current member of the group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-MOVE GRPMEMber--member_name--from_group--to_group-----<<
```

Parameters

- member_name (Required)
Specifies the member (a server or a server group) to move.
- from_group (Required)
Specifies the server group with which the member is currently associated.
- to_group (Required)
Specifies the new server group for the member.

Example: Move a server to another server group

Move member PAYSON from REGION1 group to REGION2 group.

```
move grpmember payson region1 region2
```

Related commands

Table 1. Commands related to MOVE GRPMEMBER

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

MOVE MEDIA (Move sequential-access storage pool media)

Use this command to manage overflow storage pools. The database tracks media that is moved by using this command.

This command applies to sequential-access primary and copy storage pool volumes that are managed by an automated library (including an external library). The library does not have to be full. One or more sequential-access storage pool volumes can be processed at the same time.

Use the DAYS parameter to identify eligible volumes to be moved. Use the OVERFLOW LOCATION parameter to record the storage location for the moved media.

This command generates a background process that you can view by using the QUERY PROCESS command. To cancel, issue the CANCEL PROCESS command.

To determine whether the command was successful, issue the QUERY ACTLOG command or use the server console.

The volumes that are moved by the MOVE DRMEDIA command for offsite recovery are not processed by the MOVE MEDIA command.

The MOVE MEDIA command does not process copy storage pool volumes with a DRM STATUS value of NOTMOUNTABLE, COURIER, or VAULT.

Privilege class

To issue this command, you must have one of the following privilege classes:

- If the CMD parameter is NOT specified: operator or system privilege.
- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO: operator, unrestricted storage, or system privilege.
- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES (the default): system privilege.

Syntax

```
>>-MOVE MEDia--volume_name--STGpool-----pool_name----->
    .-Days-----0-----.
>--+-----+----->
    '-Days-----days-'
>--+-----+----->
    '-WHEREState-----+MOUNTABLEInlib-----+'
                          '-MOUNTABLENotinlib-'
>--+-----+----->
    |               .-,-----|
    |               V          ||
    '-WHERESTATUSs-----+FULL-----+--+'
                          +-FILLing-+
                          '-EMPTy---'
>--+-----+-----+----->
    '-ACCess-----+READWrite-+-'   '-OVFLocation-----location-'
                          '-READOnly--'
    .-REMove-----Bulk-----.
>--+-----+-----+----->
    '-REMove-----+No-----+'   '-CMd-----"command"- '
                          +-Yes--+
                          '-Bulk-'
    .-APPend-----No-----.
>--+-----+-----+----->
    '-CMDFilename-----file_name-'   '-APPend-----+No--+-'
                                          '-Yes-'
    .-CHECKLabel-----Yes-----.
>--+-----+-----+-----><
    '-CHECKLabel-----+Yes-+-'   '-CAP-----x,y,z---'
                          '-No--'
```

Parameters

volume_name (Required)

Specifies the name of the sequential access primary or copy storage pool volume to be processed. You can use a wildcard character to specify the name. All matching volumes are considered for processing.

STGpool (Required)

Specifies the name of the sequential access primary or copy storage pool that is used to select the volumes for processing. You can use a wildcard character to specify the name. All matching storage pools are processed. If the storage pool specified is not managed by an automated library, no volumes are processed.

Days

Specifies the number of days that must elapse after the volume is written or read before the volume is eligible for processing by the command. This parameter is optional. You can specify a number from 0 to 9999. The default value is 0. The most recent of the volumes' last written date or last read date is used to calculate the number of days elapsed.

WHEREState

Specifies the current state of the volumes to be processed. This parameter is used to restrict processing to the volumes that are in the specified state. This parameter is optional. The default value is MOUNTABLEINLIB.

Possible values are:

MOUNTABLEInlib

Specifies that storage pool volumes are to move from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB state. Volumes in the MOUNTABLEINLIB state contain valid data and are in the library.

MOUNTABLENotinlib

Specifies that storage pool volumes are to change from the MOUNTABLENOTINLIB state back to the MOUNTABLEINLIB state. Volumes in the MOUNTABLENOTINLIB state might contain valid data and are in the overflow location.

- For empty scratch volumes, the MOVE MEDIA command deletes the volume records so that they can be used again.
- For private volumes, the MOVE MEDIA command resets the volume location to blank, changes the volumes' state to CHECKIN, and changes the last update date to the current date.
- For scratch volumes with data, the MOVE MEDIA command resets the volume location to blank, changes the volumes' state to CHECKIN, and changes the last update date to the current date.

Attention: Volumes in the CHECKIN state might contain valid data and must be checked into the library.

WHERESTATUS

Specifies that the move process must be restricted by volume status. This parameter is optional. You can specify more than one status in a list by separating each status with a comma and no intervening spaces. If you do not specify this parameter, volumes moved from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB state are restricted to only full volumes, and volumes moved from the MOUNTABLENOTINLIB state to the MOUNTABLEINLIB state are restricted to only empty volumes.

Possible values are:

FULL

Moves volumes with a status of FULL.

FILLing

Moves volumes with a status of FILLING.

EMPTy

Moves volumes with a status of EMPTY.

ACCess

Specifies how users and system processes access files in the storage pool volume that is moved out from an automated library and stored in an overflow location by the MOVE MEDIA command. This parameter is optional. If you do not specify this parameter, moving volumes from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB process updates the volumes' access mode to READONLY, and moving volumes from the MOUNTABLENOTINLIB state to the MOUNTABLEINLIB process updates the volumes' access mode to READWRITE.

Possible values are:

READWrite

Specifies that users and system processes can read from and write to files stored on the volume that is in the overflow location. If this value is specified, IBM Spectrum Protect™ requests the volume to be checked into the library when the volume is needed for a read or write operation.

READOnly

Specifies that users and system processes can read but not write to files that are stored on the volume that is in the overflow location. The server requests the volume to be checked into the library only when the volume is needed for a read operation.

OVFLocation

Specifies the overflow location that is the destination of the volumes that are being processed. The maximum length of the location name is 255 characters. The location name information must be enclosed in quotation marks if it contains any blank characters. If you do not specify an overflow location and the storage pool also has no overflow location identified, the server changes the location of the ejected volume to a null string ("").

REMove

Specifies that the server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values are YES, BULK, and NO. The default is BULK. The response of the server to each of those options and the default values are described in the following tables.

349X libraries: The following table shows how the server responds for 349X libraries.

Table 1. How the Server Responds for 349X Libraries

REMOVE=YES	REMOVE=BULK	REMOVE=NO
The 3494 Library Manager ejects the cartridge to the convenience I/O station.	The 3494 Library Manager ejects the cartridge to the high-capacity output facility.	The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the INSERT category for use by other applications.

SCSI libraries: The following table shows how the server responds to YES, BULK, and NO for SCSI libraries.

Table 2. How the Server Responds for SCSI Libraries

If a library...	And REMOVE=YES...	And REMOVE=BULK...	And REMOVE=NO
Does not have entry/exit ports	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
Has entry/exit ports and an entry/exit port is available	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
Has entry/exit ports, but no ports are available	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server waits for an entry/exit port to be made available.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.

ACSLs libraries: The following table shows how the server responds for ACSLS libraries.

Table 3. How the Server Responds for ACSLS Libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station. The server then deletes the volume entry from the server library inventory. While moving volumes from the MOUNTABLE state with REMOVE=YES specified, the MOVE MEDIA command uses more than one slot in the CAP for a StorageTek library with ACSLS.	The server does not eject the cartridge. The server deletes the volume entry from the server library inventory and leaves the volume in the library.

External libraries: The following table shows how the server responds for external libraries.

Table 4. How the Server Responds for External Libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station. The server then deletes the volume entry from the server library inventory.	The server does not eject the cartridge. The server deletes the volume entry from the server library inventory and leaves the volume in the library.

CMD

Specifies the creation of executable commands. This parameter is optional. You must enclose your command specification in quotation marks. The maximum length of the command specification is 255 characters. For each volume successfully processed by the MOVE MEDIA command, the server writes the associated commands to a file. Specify the file name with the CMDFILENAME parameter.

AIX Linux If you do not specify the file name, the MOVE MEDIA command generates a default file name by appending the string exec.cmds.media to the IBM Spectrum Protect server directory.

Windows If you do not specify the file name, the MOVE MEDIA command generates a default file name by appending the string exec.cmd.media to the IBM Spectrum Protect server directory.

If the length of the command that is written to the file exceeds 255 characters, it is split into multiple lines and a continuation character, +, is added to all but the last line of the command. You must alter the continuation character according to the requirements of the product that runs the commands.

If you do not specify CMD, the MOVE MEDIA command might not generate any executable commands.

string

Specifies the string to build an executable command. You can specify any free form text for the string. Enclose the full string in quotation marks. For example, the following is a valid executable command specification:

```
CMD="UPDATE VOLUME &VOL"
```

The following is an invalid executable command specification:

```
CMD=""UPDATE VOLUME" &VOL"
```

substitution

Specifies a variable for which you want the command to substitute a value. The possible substitution variables are:

&VOL

Substitute the volume name for &VOL. You can specify lowercase characters, &vol. No spaces or blanks are allowed between ampersand, &, and VOL. If there are spaces or blanks between ampersand and VOL, the MOVE MEDIA command treats them as strings and no substitution is set. If &VOL is not specified, no volume name is set in the executable command.

&LOC

Substitute the volume location for &LOC. You can specify lowercase characters, &loc. No spaces or blanks are allowed between ampersand, &, and LOC. If there are spaces or blanks between ampersand and LOC, the MOVE MEDIA command treats them as strings and no substitution is set. If &LOC is not specified, no location name is set in the executable command.

&VOLDSN

Substitute the volume file name for &VOLDSN. An example of a storage pool tape volume file name that uses the default prefix ADSM is ADSM.BFS. If &VOLDSN is not specified, no volume file name is set in the executable command.

&NL

Substitute a new line character for &NL. When &NL is specified, the MOVE MEDIA command splits the command at the position where the &NL is and does not append any continuation character. The user is responsible for specifying the correct continuation character before the &NL if one is required. The user is also responsible for the length of the line written. If the &NL is not specified and the length of the command line exceeds 255, the command line is split into multiple lines and a continuation character, +, is added to all but the last line of the command.

CMDFilename

Specifies the full path name of a file that contains the commands that are specified with CMD. This parameter is optional. The maximum length of the file name is 1279 characters.

AIX Linux If you do not specify a file name, the MOVE MEDIA command generates a default file name by appending the string `exec.cmds.media` to the IBM Spectrum Protect server directory. The server directory is the current working directory of the IBM Spectrum Protect server process.

Windows If you do not specify a file name, the MOVE MEDIA command generates a default file name by appending the string `exec.cmd.media` to the IBM Spectrum Protect server directory. The server directory is the current working directory of the IBM Spectrum Protect server process.

The MOVE MEDIA command automatically allocates the file name that is specified or generated. If the file name exists, you can use the APPEND=YES parameter to add to the file. Otherwise, the file is overwritten. If a file is accidentally overwritten and you must run the commands that were in the file, issue the QUERY MEDIA command to rebuild the executable commands for the desired volumes. If the MOVE MEDIA command fails after the command file is allocated, the file is not deleted.

APPend

Specifies to write at the beginning or ending of the command file data. The default is NO. Possible values are:

No

Specifies to write the data from the beginning of the command file. If the command file exists, its contents are overwritten.

Yes

Specifies to append the command file by writing at the end of the command file data.

CHECKLabel

Specifies whether the server reads volume labels for sequential media. For SCSI devices, you can suppress label checking by setting the CHECKLabel to NO. This parameter is not applicable to 349X libraries. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the server attempts to read the media label. Reading the media label verifies that the correct volume is being checked out.

No

Specifies that the server does not attempt to read media label. This increases performance because the read process does not occur.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (x,y,z), issue the QUERY CAP command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

x

The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.

y

The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.

z

The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Example: Move all full volumes out of the library

Move all full volumes that are in the ARCHIVE sequential primary storage pool out of the library.

```
move media * stgpool=archive
```

Example: Generate the checkin commands

Generate the CHECKIN LIBVOLUME commands for full and partially full volumes that are in the ONSITE.ARCHIVE primary storage pool and stored in the overflow location, Room 2948/Bldg31.

AIX | **Linux** MOVE MEDIA creates the executable commands in /tsm/move/media/checkin.vols

Windows MOVE MEDIA creates the executable commands in c:\tsm\move\media\checkin.vols

```
move media * stgpool=onsite.archive
wherestate=mountablenotinlib wherestatus=full,filling
ovflocation=room2948/bldg31
cmd="checkin libvol lib3494 &vol status=private"
cmdfilename=/tsm/move/media/checkin.vols
```

```
checkin libvolume lib3494 TAPE04 status=private
checkin libvolume lib3494 TAPE13 status=private
checkin libvolume lib3494 TAPE14 status=private
```

Tip: Run the CHECKIN LIBVOLUME commands by issuing the MACRO command with the following as the macro name:

- **AIX** | **Linux** /tsm/move/media/checkin.vols
- **Windows** c:\tsm\move\media\checkin.vols

Related commands

Table 5. Commands related to MOVE MEDIA

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY MEDIA	Displays information about storage pool volumes moved by the MOVE MEDIA command.
QUERY PROCESS	Displays information about background processes.

MOVE NODEDATA (Move data by node in a sequential access storage pool)

Use this command to move data that is in a sequential-access storage pool. You can move data for one or more nodes, a group of file spaces, or for a group of collocated nodes. You can also move selected file spaces for a single node. The data can be in a primary storage pool, a copy storage pool, or an active-data pool.

This command is helpful for reducing the number of volume mounts during client restore or retrieve operations by consolidating data for a specific node within a storage pool, or to move data to another storage pool. For example, you can use this command for moving data to a random-access storage pool in preparation for client restore processing.

Ensure that the access mode of the volumes from which you are moving the node data is read/write or read-only and that the access mode of the volumes to which you are moving the node data is set to read/write. This operation will not move data on volumes with access modes of offsite, unavailable, or destroyed.

The MOVE NODEDATA command takes two forms, depending on whether you are moving data only for selected filespace. The syntax and parameters for each form are defined separately.

Restriction: You cannot move node data into or out of a storage pool that is defined with a CENTERA device class.

Table 1. Commands related to MOVE NODEDATA

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY OCCUPANCY	Displays file space information by storage pool.
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE COLLOGROUP	Updates the description of a collocation group.

- MOVE NODEDATA (Move data in file spaces for one or more nodes or a collocation group)
Use this command to move data in file spaces that belong to; one or more nodes, a node collocation group, or a file space collocation group.
- MOVE NODEDATA (Move data from selected file spaces of a single node)
Use this command to move data for selected file spaces belonging to a single node.

MOVE NODEDATA (Move data in file spaces for one or more nodes or a collocation group)

Use this command to move data in file spaces that belong to; one or more nodes, a node collocation group, or a file space collocation group.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the source storage pool. If your authorization is restricted storage privilege and you are moving data to another storage pool, you need the appropriate authority for the destination storage pool.

Syntax

```

      .-,------.
      v          |
>>-MOVE NODEdata--+---node_name+-----+----->
      '-COLLOGroup-----group_name-'

>--FROMstgpool-----source_pool_name----->

>--+-----+----->
      '-TOstgpool-----destination_pool_name-'

      .-Type-----ANY------.
>--+-----+----->
      '-Type-----+ANY-----+'
          +-Backup-----+
          +-ARchive-----+

```

```

        '-SPacemanaged-'
    .-MAXPRocess-----1----- .-Wait----No-----
>-----+-----+-----+-----+-----+-----+-----+-----+----->
    '-MAXPRocess-----num_processes-' '-Wait----+No--+-'
                                     '-Yes-'

        (1)
    .-RECONStruct-----No or Yes-----
>-----+-----+-----+-----+-----+-----+-----+-----+-----><
    '-RECONStruct-----+No--+-'
                                     '-Yes-'

```

Notes:

1. The default is NO if either the source or target storage pool is random access. The default is YES if both the source and target storage pools are sequential access.

Parameters

node_name (Required unless the COLLOGROUP parameter is specified)

Specifies the node name that is related to the data that is moved with this command. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

COLLOGroup (Required unless the node_name parameter is specified)

Specifies the name of the collocation group whose data is to be moved. Data for all nodes and file spaces that belong to the collocation group are moved.

FROMstgpool (Required)

Specifies the name of a sequential-access storage pool that contains data to be moved. This storage pool must be in the NATIVE or NONBLOCK data format.

TOstgpool

Specifies the name of a storage pool to where the data is moved. This storage pool must be in the NATIVE or NONBLOCK data format. This parameter is optional and does not apply when the source storage pool is a copy storage pool or an active-data pool. That is, if the source storage pool is a copy storage pool the destination must be the same copy storage pool. Similarly, if the source storage pool is an active-data pool, the destination must be the same active-data pool. If a value is not specified, data is moved to other volumes within the source pool.

Important: If you are moving data within the same storage pool, there must be volumes available that do not contain the node data that you are moving. That is, the server cannot use volumes that contain the data to be moved as destination volumes.

Type

Specifies the type of files to be moved. This parameter is optional. The default value is ANY. If the source storage pool is an active-data pool, the only valid values are ANY and BACKUP. However, only the active versions of backup data are moved if TYPE=ANY. Specify one of the following values:

ANY

Specifies that all types of files are moved.

Backup

Specifies that backup files are moved.

ARchive

Specifies that archive files are moved. This value is not valid for active-data pools.

SPacemanaged

Specifies that space-managed files (files that were migrated by an IBM Spectrum Protect™ for Space Management client) are moved. This value is not valid for active-data pools.

MAXPRocess

Specifies the maximum number of parallel processes to use for moving data. This parameter is optional. You can specify a value from 1 to 999, inclusive. The default value is 1. Increasing the number of parallel processes usually improves throughput.

When you determine this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect system activity. The mount points and drives also depend on the mount limits of the device classes for the sequential access storage pools that are involved in the move. Each process needs a mount point for storage pool volumes, and, if the device type is not FILE, each process also needs a drive.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Specify one of the following values:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If a background process is canceled, some files might move before the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

RECONStruct

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when you move the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

You can specify one of the following values:

No

Specifies that reconstruction of file aggregates are not run during the move.

Yes

Specifies that reconstruction of file aggregates are run during the move. You can specify only this option when both the source and the target storage pools are sequential-access.

Move a specific node's data from a tape storage pool to a disk storage pool

Move all data that belongs to node MARY that is stored in storage pool TAPEPOOL. Data can be moved to disk storage pool BACKUPPOOL.

```
move nodedata mary
  fromstgpool=tapepool tostgpool=backuppool
```

Move data for a node collocation group from one storage pool to another

Move all data for node collocation group NODEGROUP1 from storage pool SOURCEPOOL to storage pool TARGETPOOL.

```
move nodedata colloggroup=nodegroup1 fromstgpool=sourcespool tostgpool=targetpool
```

Move data for a file space collocation group from one storage pool to another

Move all data for file space collocation group FSGROUP1 from storage pool SOURCEPOOL2 to storage pool TARGETPOOL2.

```
move nodedata colloggroup=fsgroup1 fromstgpool=sourcespool2 tostgpool=targetpool2
```

MOVE NODEDATA (Move data from selected file spaces of a single node)

pool. Similarly, if the source storage pool is an active-data pool, the destination must be the same active-data pool. If a value is not specified, data is moved to other volumes within the source pool.

Important: If you are moving data within the same storage pool, there must be volumes available that do not contain the node data you are moving. That is, the server cannot use volumes that contain the data to be moved as destination volumes.

FILEspace

Specifies the name of the non-Unicode file space that contains data to be moved. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. This parameter is optional. If you do not specify a value for this parameter and values for UNIFILESPACE or the FSID or both, non-Unicode file spaces are not moved.

UNIFILESpace

Specifies the name of the Unicode file space that contains data to be moved. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. This parameter is optional. If you do not specify a value for this parameter and values for FILESPACE or the FSID or both, non-Unicode file spaces are not moved.

FSID

Specifies file space identifiers (FSIDs) for the file spaces to be moved. Separate multiple names with commas and no intervening spaces. This parameter is optional.

Type

Specifies the type of files to be moved. This parameter is optional. The default value is ANY. If the source storage pool is an active-data pool, the only valid values are ANY and BACKUP. However, only the active versions of backup data are moved if TYPE=ANY. Possible values are:

ANY

Specifies that all types of files are moved.

Backup

Specifies that backup files are moved.

ARchive

Specifies that archive files are moved. This value is not valid for active-data pools.

SPacemanaged

Specifies that space-managed files (files that were migrated by an IBM Spectrum Protect™ for Space Management client) are moved. This value is not valid for active-data pools.

MAXPRocess

Specifies the maximum number of parallel processes to use for moving data. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1. Increasing the number of parallel processes should improve throughput.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the move. Each process needs a mount point for storage pool volumes, and, if the device type is not FILE, each process also needs a drive.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If a background process is canceled, some files may have already moved before the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

RECONStruct

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that has accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when moving the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

Possible values are:

No

Specifies that reconstruction of file aggregates will not be performed during the move.

Yes

Specifies that reconstruction of file aggregates will be performed during the move. You may only specify this option when both the source and the target storage pools are sequential-access.

Example: Move a node's non-Unicode and Unicode data

Move data for node TOM in storage pool TAPEPOOL. Restrict movement of data to files in non-Unicode file spaces as well as Unicode file spaces, \\jane\d\$. Data should be moved to disk storage pool BACKUPPOOL.

```
move nodedata tom
  fromstgpool=tapepool tostgpool=backuppool
  filespace=* unifilespace=\\jane\d$
```

Example: Move all node data from tape storage pools to a disk storage pool

Move all data for node SARAH, from all primary sequential-access storage pools (for this example, TAPEPOOL*) to DISKPOOL. To obtain a list of storage pools that contain data for node SARAH, issue either of the following QUERY OCCUPANCY or SELECT commands:

```
query occupancy sarah

SELECT * from OCCUPANCY where node_name='sarah'
```

Attention: For this example assume that the results were TAPEPOOL1, TAPEPOOL4, and TAPEPOOL5.

```
move nodedata sarah
  fromstgpool=tapepool1 tostgpool=DISKPOOL

move nodedata sarah
  fromstgpool=tapepool4 tostgpool=DISKPOOL

move nodedata sarah
  fromstgpool=tapepool5 tostgpool=DISKPOOL
```

Example: Move a node's non-Unicode and Unicode file spaces

The following is an example of moving non-Unicode and Unicode file spaces for a node. For node NOAH move non-Unicode filespace \\servtuc\d\$ and Unicode file space \\tsmserv1\e\$ that has a filespace ID of 2 from sequential access storage pool TAPEPOOL to random access storage pool DISKPOOL.

```
move nodedata noah
  fromstgpool=tapepool tostgpool=diskpool
  filespace=\\tsmserv1\d$ fsid=2
```

NOTIFY SUBSCRIBERS (Notify managed servers to update profiles)

Use this command on a configuration manager to notify one or more managed servers to request that their configuration information be immediately refreshed.

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-PERForm LIBAction--library_name----->
>----ACTion---+--DEFine--| A |----->
          +-DELeTe-----+
          +-RESet--| B |---+
          '-QUIesce-----'

          .-PREView-----No-----.
>--+-----+-----+-----><
  '-SOURCe-----source_name-' '-PREView-----+--Yes--+'
                                   '-No--'

A (DEFine)

|--+-----+----->
  '-DEVIce-----library_device_name-'

  .-PREFix-----library_name-----
>--+-----+-----|
  '-PREFix-----drive_prefix_name-'

B (RESet)

          .-DRIVEsonly-----No-----
|----ACTion---+--RESet---+-----|
          '-DRIVEsonly-----+--Yes--+'
                                   '-No--'
```

Parameters

library_name (Required)

Specifies the name of the library to be defined or deleted. The maximum length of this name is 30 characters unless you are issuing PERFORM LIBACTION with ACTION=DEFINE and using the default PREFIX value. In that case, the maximum length of the name is 25 characters.

ACTion

Specifies the action for the PERFORM LIBACTION command. Possible values are:

DEFine

Specifies that drives and their paths are defined for the specified library. SAN discovery must be enabled before you specify this parameter value.

DELeTe

Specifies that drives and their paths are deleted for the specified library.

RESet

Specifies that drives and their paths are updated online for the specified library.

DRIVEsonly

Specifies that only drives are updated online for the specified library.

Possible values are:

No

Specifies that drives and paths are updated online.

Yes

Specifies that only drives are updated online.

QUIesce

Specifies that drives are updated offline.

DEVIce

Specifies the library device name that is used when you define paths if a path to the library is not already defined. If a path is already defined, the DEVICE parameter is ignored. The maximum length for this value is 64 characters. This parameter is optional.

PREFix

Specifies the prefix that is used for all drive definitions. For example, a PREFIX value of *DR* creates drives *DR0*, *DR1*, *DR2*, for as many drives as are created. If a value is not specified for the PREFIX parameter, the library name is used as the prefix for drive definitions. The maximum length for this value is 25 characters.

SOURCE

Specifies the source server name to be used when you define or delete drive path definitions on a library client or LAN-free client. Use this parameter only if the drives in the library are set up for the local server. If no value is specified for the SOURCE parameter, the local server name, which is the default, is used. The maximum length for the source name is 64 characters.

If you specify the SOURCE parameter, you can RESET only paths from specified SOURCE values. The SOURCE parameter is not compatible with the RESET DRIVESONLY=YES or QUIESCE options.

If a source name other than the local server name is specified with ACTION=DEFINE, drive path definitions are defined with the token value of UNDISCOVERED. The path definitions are then updated dynamically by library clients that support SAN Discovery the first time the drive is mounted.

PREVIEW

Specifies the output of all commands that are processed for PERFORM LIBACTION before the command is issued. The PREVIEW parameter is not compatible with the DEVICE parameter. If you are issuing the PERFORM LIBACTION command to define a library, you cannot specify both the PREVIEW and the DEVICE parameter.

Possible values are:

No

Specifies that a preview of the commands that are issued for PERFORM LIBACTION is not displayed.

Yes

Specifies that a preview of the commands that are issued for PERFORM LIBACTION is displayed.

Example: Define a shared library

Assume that you are working in a SAN and that you configured a library manager named LIBMGR1. Now, define a library that is named SHAREDTSM to a library client server named LIBCL1.

Issue DEFINE LIBRARY from the library client server, LIBCL1:

```
define library sharedtsm libtype=shared primarylibmanager=libmgr1
```

Then, issue PERFORM LIBACTION from the library manager, LIBMGR1, to define the drive paths for the library client:

```
perform libaction sharedtsm action=define source=libcl1
```

Note: The SANDISCOVERY option must be supported and enabled on the library client server.

Example: Define a library with four drives

Define a SCSI library named KONA:

```
define library kona libtype=scsi
```

Then issue the PERFORM LIBACTION command to define drives and paths for the library:

AIX

```
perform libaction kona action=define device=/dev/lb3  
prefix=dr
```

The server then runs the following commands:

```
define path server1 kona srct=server destt=library  
device=/dev/lb3  
define drive kona dr0  
define path server1 dr0 srct=server destt=drive library=kona  
device=/dev/mt1  
define drive kona dr1  
define path server1 dr1 srct=server destt=drive library=kona  
device=/dev/mt2  
define drive kona dr2  
define path server1 dr2 srct=server destt=drive library=kona  
device=/dev/mt3  
define drive kona dr3
```

```
define path server1 dr3 srct=server destt=drive library=kona
device=/dev/mt4
```

Linux

```
perform libaction kona action=define device=/dev/tmscsi/lb3
prefix=dr
```

The server then runs the following commands:

```
define path server1 kona srct=server destt=library
device=/dev/tmscsi/lb3
define drive kona dr0
define path server1 dr0 srct=server destt=drive library=kona
device=/dev/tmscsi/mt1
define drive kona dr1
define path server1 dr1 srct=server destt=drive library=kona
device=/dev/tmscsi/mt2
define drive kona dr2
define path server1 dr2 srct=server destt=drive library=kona
device=/dev/tmscsi/mt3
define drive kona dr3
define path server1 dr3 srct=server destt=drive library=kona
device=/dev/tmscsi/mt4
```

Windows

```
perform libaction kona action=define device=lb0.0.0.2
prefix=dr
```

The server then runs the following commands:

```
define path server1 kona srct=server destt=library
device=lb0.0.0.2
define drive kona dr0
define path server1 dr0 srct=server destt=drive library=kona
device=mt0.1.0.2
define drive kona dr1
define path server1 dr1 srct=server destt=drive library=kona
device=mt0.2.0.2
define drive kona dr2
define path server1 dr2 srct=server destt=drive library=kona
device=mt0.3.0.2
define drive kona dr3
define path server1 dr3 srct=server destt=drive library=kona
device=mt0.4.0.2
```

Related commands

Table 1. Commands related to PERFORM LIBACTION

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.

Command	Description
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE PATH	Changes the attributes associated with a path.

PING SERVER (Test the connection between servers)

Use this command to test the connection between the local server and a remote server.

Important: The name and password of the administrator client issuing this command must also be defined on the remote server. If the remote server is at the current level, the server credentials are verified automatically when you run the PING SERVER command. If the remote server is not at the current level, the server credentials are not verified.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-PING SERVER--server_name-----<<
```

Parameters

server_name (Required)
Specifies the name of the remote server.

Example: Ping a server

Test the connection to server FRED.

```
ping server fred
```

Related commands

Table 1. Commands related to PING SERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
QUERY SERVER	Displays information about servers.

PREPARE (Create a recovery plan file)

Use this command to create a recovery plan file, which contains the information that is needed to recover an IBM Spectrum Protect™ server. You can store a recovery plan file on a file system that is accessible to the source server or on a target server.

You can use the QUERY ACTLOG command to view whether the PREPARE command was successful.

You can also view this information from the server console or, if the WAIT parameter equals YES, an administrative client session.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
.-Source-----DBBackup-----.
```

```

>>-Prepare-----+-----+-----+----->
      '-Source-----+DBBackup-----+'
                '-DBSnapshot-'

>-----+-----+-----+----->
      '-DEVclass-----device_class_name-'

>-----+-----+-----+----->
      '-PLANPrefix-----prefix-' '-INSTRPrefix-----prefix-'

>-----+-----+-----+----->
      |               .-,----- . |
      |               V             | |
      '-COPYstgpool-----pool_name+--'

>-----+-----+-----+----->
      |               .-,----- . |
      |               V             | |
      '-ACTIVEDatastgpool-----pool_name+--'

      |               .-,----- . |
      |               V             | |
      '-PRIMstgpool-----pool_name+--'

      |               .-,----- . |
      |               V             | |
      |               '-Wait-----No-----'
      |               '-Wait-----+No--+-'
      |               '-Yes-'

      |               .-,----- . |
      |               V             | |
      '-PRIMstgpool-----pool_name+--'

```

Parameters

Source

Specifies the type of database backup series that IBM Spectrum Protect assumes when generating the recovery plan file. This parameter is optional. The default is DBBACKUP. The choices are:

DBBackup

Specifies that IBM Spectrum Protect assumes the latest full database backup series.

DBSnapshot

Specifies that IBM Spectrum Protect assumes the latest database snapshot backup series.

DEVclass

Specifies the device class name that is used to create a recovery plan file object on a target server. The device class must have a device type of SERVER.

Important: The maximum capacity for the device class must be larger than the size of the recovery plan file. If the size of the recovery plan file exceeds the maximum capacity, the command fails.

The naming convention for the archive object that contains the recovery plan file on the target server is:

- **Filespace name:**
 - ADSM.SERVER
- **High-level qualifier:**
 - **AIX** | **Linux** devclassprefix/servername.yyyymmdd.hhmmss
 - **Windows** devclassprefix\servername.yyyymmdd.hhmmss
- **Low-level qualifier:**
 - RPF.OBJ.1

The recovery plan file virtual volume name as recorded in the volume history table on the source server is in the format servername.yyyymmdd.hhmmss.

If the DEVCLASS parameter is not specified, the recovery plan file is written to a file based on the plan prefix.

If SOURCE=DBBACKUP is specified or is defaulted to, the volume history entry for the recovery plan file object specifies a volume type of RPFIL. If SOURCE=DBSNAPSHOT is specified, the volume history entry specifies a volume type of RPFNSNAPSHOT.

PLANPrefix

Specifies the path name prefix that is used in the recovery plan file name. This parameter is optional.

- **AIX** | **Linux** The maximum length is 250 characters.
- **Windows** The maximum length is 200 characters.

Windows Specifies the path name prefix that is used in the recovery plan file name.

IBM Spectrum Protect appends to the prefix the sortable date and time format `yyyymmdd.hhmmss`. For example: 20081115.051421.

AIX | **Linux** The prefix can be one of the following:

Directory path

End the prefix with the forward slash (/). For example:

```
PLANPREFIX=/admsrv/recplans/
```

The resulting file name would look like this:

```
/admsrv/recplans/20081115.051421
```

Directory path followed by a string

IBM Spectrum Protect treats the string as part of the file name. For example:

```
PLANPREFIX=/admsrv/recplans/accounting
```

The resulting file name looks like this:

```
/admsrv/recplans/accounting.20081115.051421
```

Note the period before the date and time.

String only

IBM Spectrum Protect specifies the directory path. IBM Spectrum Protect uses the name of the current working directory. For example, the current working directory is `/opt/tivoli/tsm/server/bin` and you specify the following parameter:

```
PLANPREFIX=shipping
```

The resulting file name looks like this:

```
/opt/tivoli/tsm/server/bin/shipping.20081115.051421
```

Note the period before the date and time.

Windows The prefix can be one of the following:

Directory path

End the prefix with the back slash (\). For example:

```
PLANPREFIX=c:\admsrv\recplans\
```

The resulting file name looks like this:

```
c:\admsrv\recplans\20081115.051421
```

Tip: If you issue the PREPARE command from the administrative command line client and the last character in the command line is a back slash, it is interpreted as a continuation character. To avoid this, place the prefix value in double quotation marks. For example:

```
PLANPREFIX="c:\admsrv\recplans\"
```

Directory path followed by a string

IBM Spectrum Protect treats the string as part of the file name. For example:

```
PLANPREFIX=c:\admsrv\recplans\accounting
```

The resulting file name looks like this:

```
c:\admsrv\recplans\accounting.20081115.051421
```

Note the period before the date and time.

String only

IBM Spectrum Protect appends the date and time in the `yyyymmdd.hhmmss` format (note the period before the date and time) to the prefix. The directory path used by the PREPARE command is the directory representing this "instance" of the IBM Spectrum Protect server. Typically, this directory is the original IBM Spectrum Protect server

installation directory. For example, the directory representing this instance of the server is c:\Program Files\Tivoli\TSM;\server2 , and you issue a PREPARE command with the following parameter:

```
PLANPREFIX=shipping
```

The resulting recovery plan filename is:

```
c:\Program Files\Tivoli\TSM;\server2\shipping.20081115.051421
```

If the PLANPREFIX parameter is not specified, IBM Spectrum Protect selects the prefix in one of these ways:

- If the SET DRMPREFIX command has been issued, IBM Spectrum Protect uses the prefix specified in that command.
- **Windows** If the SET DRMPREFIX command is not defined, IBM Spectrum Protect uses as the path the directory representing this “instance” of the IBM Spectrum Protect server, which is typically the original IBM Spectrum Protect server installation directory. For example, the directory representing this instance of the server is the following:

```
c:\Program Files\Tivoli\TSM;\server2
```

The resulting recovery plan file name is the following:

```
c:\Program Files\Tivoli\TSM;\server2\20081115.051421
```

- **AIX** | **Linux** If the SET DRMPREFIX command has not been issued, IBM Spectrum Protect uses the directory path name of the current working directory. For example, the current working directory is the following:

```
/opt/tivoli/tsm/server/bin
```

The resulting file name looks like this:

```
/opt/tivoli/txm/server/bin/20081115.051421
```

INSTRPrefix

Specifies the prefix of the path name used by IBM Spectrum Protect to locate the files that contain the recovery instructions. The maximum length is **AIX** | **Linux** 250 **Windows** 200 characters.

AIX | **Linux** The prefix can be one of the following:

Directory path

End the prefix with the forward slash (/). For example:

```
INSTRPREFIX=/admsrv/recinstr/  
/admsrv/recinstr/RECOVERY.INSTRUCTIONS.GENERAL
```

Directory path followed by a string

IBM Spectrum Protect treats the string as part of the file name. For example:

```
INSTRPREFIX=/admsrv/recinstr/accounts
```

IBM Spectrum Protect appends the appropriate recovery plan file stanza name. For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name is:

```
/admsrv/recinstr/accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

String only

- IBM Spectrum Protect specifies the directory path and appends the appropriate recovery plan file stanza name. IBM Spectrum Protect uses the name of the current working directory. For example, the current working directory is /opt/tivoli/tsm/server/bin and you specify the following parameter:

```
INSTRPREFIX=shipping
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name looks like this:

```
/opt/tivoli/tsm/server/bin/shipping.RECOVERY.INSTRUCTIONS.GENERAL
```

Windows The prefix can be one of the following:

Directory path

End the prefix with the back slash (\). For example:


```
INSTRPREFIX=c:\admsrv\recinstr\
```

IBM Spectrum Protect appends the appropriate recovery plan file stanza name. For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name is:

```
c:\admsrv\recinstr\RECOVERY.INSTRUCTIONS.GENERAL
```

Tip: If you issue the PREPARE command from the administrative command line client and the last character in the command line is a back slash, it is interpreted as a continuation character. To avoid this, place the prefix value in double quotation marks. For example:

```
INSTRPREFIX="c:\admsrv\recinstr\"
```

Directory path followed by a string

IBM Spectrum Protect treats the string as part of the file name. For example:

```
INSTRPREFIX=c:\admsrv\recinstr\accounts
```

IBM Spectrum Protect appends the appropriate recovery plan file stanza name. For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name is:

```
c:\admsrv\recinstr\accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

String only

IBM Spectrum Protect specifies the directory path and appends the appropriate recovery plan file stanza name. IBM Spectrum Protect appends the recovery plan file stanza name to the prefix. If the prefix is only a string, the directory path used by the PREPARE command is the directory representing this instance of the IBM Spectrum Protect server. This is typically the original IBM Spectrum Protect server installation directory. For example, the directory representing this instance of the server is c:\Program Files\Tivoli\TSM;\server2, and you issue a PREPARE command with the following parameter:

```
INSTRPREFIX=dock
```

The resulting recovery plan filename is:

```
c:\Program Files\Tivoli\TSM;\server2\shipping.20081115.051421
```

If you do not specify the INSTRPREFIX parameter, IBM Spectrum Protect selects the prefix in one of these ways:

- If the SET DRMINSTRPREFIX command has been issued, IBM Spectrum Protect uses the prefix specified in that command.
- **Windows** If the SET DRMINSTRPREFIX command has not been issued, IBM Spectrum Protect uses as the path the directory representing this “instance” of the IBM Spectrum Protect server, which is typically the original server installation directory. For example, the directory representing this instance of the server is the following:

```
c:\Program Files\Tivoli\TSM;\server2
```

The resulting recovery plan file name is the following:

```
c:\Program Files\Tivoli\TSM;\server2\RECOVERY.INSTRUCTIONS.GENERAL
```

- **AIX** | **Linux** If the SET DRMINSTRPREFIX command has not been issued, IBM Spectrum Protect uses the current working directory. For example, if the current working directory is /opt/tivoli/tsm/server/bin, for the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
/opt/tivoli/tsm/server/bin/RECOVERY.INSTRUCTIONS.GENERAL
```

PRIMstgpool

Specifies the names of the primary storage pools that you want to restore. Separate the storage pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, IBM Spectrum Protect selects the storage pools as follows:

- If the SET DRMPRIMSTGPOOL command has been issued, IBM Spectrum Protect includes the primary storage pools named in that command.
- If the SET DRMPRIMSTGPOOL command has not been issued, IBM Spectrum Protect includes all the primary storage pools.

COPYstgpool

Specifies the names of the copy storage pools used to back up the primary storage pools that you want to restore (see the PRIMSTGPOOL parameter). Separate storage pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, IBM Spectrum Protect selects the storage pools as follows:

- If the SET DRMCOPYSTGPOOL command has been issued, IBM Spectrum Protect includes those copy storage pools.
- If the SET DRMCOPYSTGPOOL command has not been issued, IBM Spectrum Protect includes all copy storage pools.

ACTIVEDatastgpool

Specifies the names of the active-data storage pools that you want to have available for offsite access. Separate active-data storage-pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, IBM Spectrum Protect selects the storage pools as follows:

- If the SET ACTIVE DATASTGPOOL command has been previously issued with valid active-data storage pool names, IBM Spectrum Protect processes those storage pools.
- If the SET ACTIVE DATASTGPOOL command has not been issued, or all of the active-data storage pools have been removed using the SET ACTIVE DATASTGPOOL command, IBM Spectrum Protect processes only the active-data pool volumes that were marked on-site at the time the PREPARE command is run. IBM Spectrum Protect will mark these volumes as UNAVAILABLE.

Wait

Specifies whether this command is processed in the background or foreground.

No

Specifies background processing. This is the default.

Yes

Specifies foreground processing.

AIX | **Linux** You cannot specify YES from the server console.

Example: Create a recovery plan file

Issue the PREPARE command and query the activity log to check the results.

```
prepare
query actlog search=prepare
```

AIX | **Linux**

```
05/03/2008 12:01:13 ANR0984I Process 3 for PREPARE started in the
BACKGROUND at 12:01:13.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
/home/guest/drmtest/prepare/tserver/DSM1509/
RECOVERY.INSTRUCTIONS.DATABASE not found.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
/home/guest/drmtest/prepare/tserver/DSM1509/
RECOVERY.INSTRUCTIONS.STGPOOL not found.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEEP.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEPSM.
05/03/2008 12:01:14 ANR6920W PREPARE: Generated replacement volume
name BACK4X@ is not valid for device type
8MM. Original volume name: BACK4X. Stanza is
PRIMARY.VOLUMES.REPLACEMENT macro.
05/03/2008 12:01:14 ANR6900I PREPARE: The recovery plan file
/home/guest/drmtest/prepare/plandir/DSM1509/
r.p.20080503.120113 was created.
05/03/2008 12:01:14 ANR0985I Process 3 for PREPARE running in the
BACKGROUND completed with completion state
SUCCESS at 12:01:14.
```

Windows

```
05/03/2008 12:01:13 ANR0984I Process 3 for PREPARE started in the
BACKGROUND at 12:01:13.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
c:\drmtest\prepare\RECOVERY.INSTRUCTIONS.DATABASE
not found.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
c:\drmtest\prepare\RECOVERY.INSTRUCTIONS.STGPOOL
```

```

not found.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEEP.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEPSM.
05/03/2008 12:01:14 ANR6920W PREPARE: Generated replacement volume
name BACK4X@ is not valid for device class 8MM.
Original volume name: BACK4X. Stanza is
PRIMARY.VOLUMES.REPLACEMENT macro.
05/03/2008 12:01:14 ANR6900I PREPARE: The recovery plan file
c:\drmtest\prepare\r.p.20080503.120113
was created.
05/03/2008 12:01:14 ANR0985I Process 3 for PREPARE running in the
BACKGROUND completed with completion state
SUCCESS at 12:01:14.

```

Related commands

Table 1. Commands related to PREPARE

Command	Description
CANCEL PROCESS	Cancels a background server process.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.
QUERY RPFFILE	Displays information about recovery plan files.
QUERY SERVER	Displays information about servers.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMINSTRPREFIX	Specifies the prefix portion of the path name for the recovery plan instructions.
SET DRMPPLANVPOSTFIX	Specifies the replacement volume names in the recovery plan file.
SET DRMPPLANPREFIX	Specifies the prefix portion of the path name for the recovery plan.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SET DRMRPFEXPIREDDAYS	Set criteria for recovery plan file expiration.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.

AIX Linux Windows

PROTECT STGPOOL (Protect data that belongs to a storage pool)

Use this command to protect data in a directory-container storage pool by storing a copy of the data in another storage pool on a replication target server or on the same server by protecting the data to tape. When you protect the directory-container storage pool, you can later try to repair damage in the storage pool by using the REPAIR STGPOOL command.

When you issue the PROTECT STGPOOL command for a directory-container storage pool, data that is stored in that storage pool is backed up to the target that you specify. The data can be backed up to the following target types:

- A directory-container storage pool on the target replication server.
Prerequisite: For the storage pool that is being protected, you must specify the target pool by using the PROTECTSTGPOOL parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.

When you regularly use the PROTECT STGPOOL command, you can typically reduce the processing time for the REPLICATE NODE command. The data extents that are already copied to the target replication server by storage pool protection operations are skipped when node replication is started.

As part of the PROTECT STGPOOL operation, processes might run to repair damaged extents in the target server's storage pool. The repair operation occurs under the following conditions:

- o Both the source server and the target server must be at V7.1.5 or later.
- o Extents that are already marked as damaged on the target server are repaired. The repair process does not run an audit process to identify damage.
- o Only target extents that match source extents are repaired. Target extents that are damaged but have no match on the source server are not repaired.

Limitations: The repair operation that runs as part of the PROTECT STGPOOL operation has the following limitations:

- o Extents that belong to objects that were encrypted are not repaired.
- o The timing of the occurrence of damage on the target storage pool and the sequence of REPLICATE NODE and PROTECT STGPOOL commands can affect whether the repair process is successful. Some extents that were stored in the target storage pool by a REPLICATE NODE command might not be repaired.

- Container-copy storage pools on the same server, protected to tape.

Prerequisite: For the storage pool that is being protected, you must specify the target storage pool by using the PROTECTLOCALSTGPOOLS parameter. For details about the parameter, see the commands for defining and updating directory-container storage pools (DEFINE STGPOOL and UPDATE STGPOOL commands).

As part of the PROTECT STGPOOL operation, volumes in the target pool might be reclaimed. The value of the RECLAIM parameter for the container-copy storage pool affects whether volumes are reclaimed. For details about the parameter, see the commands for defining and updating container-copy storage pools (DEFINE STGPOOL and UPDATE STGPOOL commands).

Restriction: You cannot schedule multiple PROTECT STGPOOL operations to run concurrently. Wait for one PROTECT STGPOOL operation to finish before you start another.

Privilege class

To issue this command, you must have system privilege.

Syntax when the target is the replication server

```

>>-PROTECT STGPOOL--source_stgpool-----+----->
                                     .-Type-----Replserver-.
                                     '-Type-----Replserver-'

    .-FORCEREconcile-----No----->
>--+-----+----->
    '-FORCEREconcile-----+No--+-'
                                     '-Yes-'

                                     (1)
    .-MAXSESSions-----10----->
>--+-----+----->
    '-MAXSESSions-----number_sessions--'

    .-Preview-----No----->
>--+-----+-----+----->
    '-Preview-----+No--+-'   '-PURGEdata-----+No-----+'
                                     '+-All-----+'
                                     '-Deleted-'

    .-Wait-----No----->
>--+-----+-----+----->>
    '-Wait-----+No--+-'   |                                     (2) |
                                     '-TRANSFERMethod-----+Tcpi-----+'
                                     '-Fasp--'

```

Notes:

1. **Linux** If the TRANSFERMETHOD parameter is set to the default value of TCPIP, the default value of the MAXSESSIONS parameter is 10. If the TRANSFERMETHOD parameter is set to FASP, the default value of the MAXSESSIONS parameter is

When you set a value for the MAXSESSIONS parameter, ensure that the available bandwidth and the processor capacity of the source and target servers are sufficient.

Tips:

- If you issue a QUERY SESSION command, the total number of sessions might exceed the number of data sessions. The difference is because of short control sessions that are used to query and set up operations.
- The number of sessions that are used for protection depends on the amount of data that is backed up. If you are backing up only a small amount of data, increasing the number of sessions provides no benefit.

Preview

Specifies whether to preview data. This parameter is optional. The default value is NO. Specify one of the following values:

No

Specifies that the data is backed up to the target server but that the data is not previewed.

Yes

Specifies that data is previewed but not backed up.

PURGEdata

Specifies that data extents are deleted from the target server. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that data extents that were deleted from the source server are deleted from the target server. New data extents are sent from the source server.

All

Specifies that all data extents are deleted from the target server, except for data extents that are referenced by other data in the target storage pool.

Deleted

Specifies that data extents that were deleted from the source server are deleted from the target server. No new data extents are sent from the source server.

RECLaim

Specifies whether reclamation runs when the PROTECT STGPOOL command is processed. Reclamation runs on the local container-copy storage pool that is the target for the protection operation. This parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that reclamation runs when the command is issued, along with the storage pool protection operation. Reclamation runs to completion, with no limitation on the number of volumes in the storage pool that are processed for reclamation.

No

Specifies that reclamation is not run when the command is issued. Only the storage pool protection operation runs.

Only

Specifies that reclamation is the only operation that runs when the command is issued. The storage pool protection operation does not run, so data in the directory-container storage pool that was updated since the last protection operation is not protected. Reclamation runs to completion, with no limitation on the number of volumes in the storage pool that are processed for reclamation.

YESLIMited

Specifies that reclamation runs when the command is issued, along with the storage pool protection operation. Reclamation runs until it reaches the reclaim limit that is defined for the container-copy storage pool. The reclaim limit is defined with the RECLAIMLIMIT parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.

ONLYLIMited

Specifies that reclamation is the only operation that runs when the command is issued. The storage pool protection operation does not run, so data in the directory-container storage pool that was updated since the last protection operation is not protected. Reclamation runs until it reaches the reclaim limit that is defined for the container-copy storage pool. The reclaim limit is defined with the RECLAIMLIMIT parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.

Wait

Specifies whether to wait for the server to process this command in the foreground. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the command is processed in the background. To monitor the background processes of this command, issue the QUERY PROCESS command.

Yes

Specifies that the command is processed in the foreground. Messages are not displayed until the command completes processing.

Restriction: You cannot specify WAIT=YES from the server console.

Linux TRANSFERMethod

Linux Specifies the method that is used for server-to-server data transfer. This parameter is optional. You can specify one of the following values:

Tcpip

Specifies that TCP/IP is used to transfer data. This value is the default.

Fasp

Specifies that Aspera® Fast Adaptive Secure Protocol (FASP®) technology is used to transfer data. Aspera FASP technology can help you optimize data transfer in a wide area network (WAN). If you specify TRANSFERMETHOD=FASP, you override any TRANSFERMETHOD parameters that you specified on the DEFINE SERVER or UPDATE SERVER commands.

Restrictions:

- Before you enable Aspera FASP technology, determine whether the technology is appropriate for your system environment and install the appropriate licenses. For instructions, see Determining whether Aspera FASP technology can optimize data transfer in your system environment. If the licenses are missing or expired, operations to protect storage pools fail.
- If WAN performance meets your business needs, do not enable Aspera FASP technology.

Example: Delete all data extents from the target server

Delete all data extents in a directory-container storage pool on the target server. The directory-container storage pool that is named POOL1 on the source server is no longer protected by the directory-container storage pool on the target server. You might delete all extents to clean the directory-container storage pool on the target server that no longer protects the source server.

```
protect stgpool pool1 purgedata=all
```

Example: Protect a storage pool and specify a maximum number of data sessions

Protect a storage pool that is named SPOOL1 on the source server by backing up the data to a target replication server, TPOOL1. Specify a maximum of 20 data sessions.

```
update stgpool spool1 protectstgpool=tpool1  
protect stgpool spool1 maxsessions=20
```

Example: Copy the storage pool data to tape

Protect a directory-container storage pool by copying the data to a container-copy storage pool on the same server. In this example, the directory-container storage pool is named SPOOL1 and the container-copy storage pool, which uses tape for storage, is named TAPES1.

1. Update the directory-container storage pool to add TAPES1 as the local storage pool for protection. The TAPES1 storage pool must be a container-copy storage pool. Issue the following command:

```
update stgpool spool1 protectlocalstgpools=tapes1
```

2. Protect the data in the directory-container storage pool with a local copy by issuing the following command:

```
protect stgpool type=local spool1
```

The data is copied to the TAPES1 storage pool.

Example: Reclaim space on tape volumes before you protect a storage pool

Reclaim space on the tape volumes that are used to protect a directory-container storage pool. Then, protect the data in the directory-container storage pool. In this example, the directory-container storage pool is named SPOOL1.

1. Reclaim space in the local container-copy storage pool that is defined as the target protection pool for SPOOL1.

```
protect stgpool spool1 type=local reclaim=only
```

2. Protect the data in the directory-container storage pool that is named SPOOL1 without running reclamation.

```
protect stgpool spool1 type=local reclaim=no
```

Table 1. Commands related to PROTECT STGPOOL

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE STGPOOL (container-copy)	Define a container-copy storage pool that stores copies of data from a directory-container storage pool.
DEFINE STGPOOL (directory-container)	Define a directory-container storage pool.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
REPAIR STGPOOL	Repairs a directory-container storage pool.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET REPLSERVER	Specifies a target replication server.
UPDATE STGPOOL (container-copy)	Update a container-copy storage pool that stores copies of data from a directory-container storage pool.

QUERY commands

Use the QUERY commands to request or display information about IBM Spectrum Protect™ objects.

- QUERY ACTLOG (Query the activity log)
- QUERY ADMIN (Display administrator information)
- QUERY ALERTTRIGGER (Query the list of defined alert triggers)
- QUERY ALERTSTATUS (Query the status of an alert)
- QUERY ASSOCIATION (Query client node associations with a schedule)
- QUERY AUDITOCUPANCY (Query client node storage utilization)
- QUERY BACKUPSET (Query a backup set)
- QUERY BACKUPSETCONTENTS (Query contents of a backup set)
- **AIX** **Linux** **Windows** QUERY CLEANUP (Query the cleanup that is required in a source storage pool)
- QUERY CLOPTSET (Query a client option set)
- QUERY COLLOGGROUP (Query a collocation group)
- QUERY CONTENT (Query the contents of a storage pool volume)
- **AIX** **Linux** **Windows** QUERY CONTAINER (Query a container)
- **AIX** **Linux** **Windows** QUERY CONVERSION (Query conversion status of a storage pool)
- QUERY COPYGROUP (Query copy groups)
- QUERY DATAMOVER (Display data mover definitions)
- **AIX** **Linux** **Windows** QUERY DAMAGED (Query damaged data in a directory-container or cloud-container storage pool)
- QUERY DB (Display database information)
- QUERY DBSPACE (Display database storage space)
- **AIX** **Linux** **Windows** QUERY DEDUPSTATS (Query data deduplication statistics)
- QUERY DEVCLASS (Display information on one or more device classes)
- QUERY DIRSPACE (Query storage utilization of FILE directories)
- QUERY DOMAIN (Query a policy domain)
- QUERY DRIVE (Query information about a drive)
- QUERY DRMEDIA (Query disaster recovery media)
- QUERY DRMSTATUS (Query disaster recovery manager system parameters)
- QUERY ENABLED (Query enabled events)
- QUERY EVENT (Query scheduled and completed events)
- QUERY EVENTRULES (Query rules for server or client events)
- QUERY EVENTSERVER (Query the event server)
- QUERY EXPORT (Query for active or suspended export operations)
- **AIX** **Linux** **Windows** QUERY EXTENTUPDATES (Query updated data extents)
- QUERY FILESPACE (Query one or more file spaces)

- QUERY LIBRARY (Query a library)
- QUERY LIBVOLUME (Query a library volume)
- QUERY LICENSE (Display license information)
- QUERY LOG (Display information about the recovery log)
- QUERY MACHINE (Query machine information)
- QUERY MEDIA (Query sequential-access storage pool media)
- QUERY MGMTCLASS (Query a management class)
- QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)
- QUERY MONITORSTATUS (Query the monitoring status)
- QUERY MOUNT (Display information on mounted sequential access volumes)
- | | | |
|-----|-------|---------|
| AIX | Linux | Windows |
|-----|-------|---------|

 QUERY NASBACKUP (Query NAS backup images)
- QUERY NODE (Query nodes)
- QUERY NODEDATA (Query client data in volumes)
- QUERY NODEGROUP (Query a node group)
- QUERY OCCUPANCY (Query client file spaces in storage pools)
- QUERY OPTION (Query server options)
- QUERY PATH (Display a path definition)
- QUERY POLICYSET (Query a policy set)
- QUERY PROCESS (Query one or more server processes)
- QUERY PROFILE (Query a profile)
- QUERY PROTECTSTATUS (Query the status of storage pool protection)
- QUERY PROXYNODE (Query proxy authority for a client node)
- QUERY PVUESTIMATE (Display processor value unit estimate)
- QUERY RECOVERYMEDIA (Query recovery media)
- QUERY REPLICATION (Query node replication processes)
- QUERY REPLNODE (Display information about replication status for a client node)
- QUERY REPLRULE (Query replication rules)
- QUERY REPLSERVER (Query a replication server)
- QUERY REQUEST (Query one or more pending mount requests)
- QUERY RESTORE (Query restartable restore sessions)
- QUERY RPFCONTENT (Query recovery plan file contents stored on a target server)
- QUERY RPFFILE (Query recovery plan file information stored on a target server)
- | | | |
|-----|-------|---------|
| AIX | Linux | Windows |
|-----|-------|---------|

 QUERY SAN (Query the devices on the SAN)
- QUERY SCHEDULE (Query schedules)
- QUERY SCRIPT (Query IBM Spectrum Protect scripts)
- QUERY SERVER (Query a server)
- QUERY SERVERGROUP (Query a server group)
- QUERY SESSION (Query client sessions)
- QUERY SHREDSTATUS (Query shredding status)
- QUERY SPACETRIGGER (Query the space triggers)
- QUERY STATUS (Query system parameters)
- QUERY STATUSTHRESHOLD (Query status monitoring thresholds)
- QUERY STGRULE (Display storage rule information)
- QUERY STGPOOL (Query storage pools)
- | | | |
|-----|-------|---------|
| AIX | Linux | Windows |
|-----|-------|---------|

 QUERY STGPOOLDIRECTORY (Query a storage pool directory)
- QUERY SUBSCRIBER (Display subscriber information)
- QUERY SUBSCRIPTION (Display subscription information)
- QUERY SYSTEM (Query the system configuration and capacity)
- | | | |
|-----|-------|---------|
| AIX | Linux | Windows |
|-----|-------|---------|

 QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)
- | | | |
|-----|-------|---------|
| AIX | Linux | Windows |
|-----|-------|---------|

 QUERY TOC (Display table of contents for a backup image)
- QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)
- QUERY VOLHISTORY (Display sequential volume history information)
- QUERY VOLUME (Query storage pool volumes)

QUERY ACTLOG (Query the activity log)

Use this command to display messages generated by the server and client. This command provides filtering options that can be used to limit the number of messages displayed and the time that it takes to process this query. If you do not specify any parameters with this command, all messages generated in the previous hour are displayed.

The activity log contains all messages that are sent to the server console under normal operation. The results of commands entered at the server console are not recorded in the activity log unless the command affects or starts a background process or client session. Error messages are displayed in the activity log.

Restriction: You cannot schedule the QUERY ACTLOG command by using the DEFINE SCHEDULE command.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-BEGINDate---current_date-.
>>-Query Actlog----->
      '-BEGINDate---date-----'

      .-BEGINTime---currenttime_minus_1_hour-.
>----->
      '-BEGINTime---time-----'

      .-ENDDate---current_date-.  .-ENDTime---current_time-.
>----->
      '-ENDDate---date-----'  '-ENDTime---time-----'

>----->
      '-MSGno---message_number-'  '-Search---string-'

>----->
      '-NODEname---node_name-'

      .-ORiginator---ALL-----
>----->
      '-ORiginator---+ALL-----+'
          +-Server-----+
          '-CLient--| A |-'

A

|----->
      '-OWNErname---owner_name-'

>----->
      '-SCHedname---schedule_name-'

>----->
      '-DOWmainname---domain_name-'

>-----|
      '-SESsnum---session_number-'

```

Parameters

BEGINDate

Specifies the beginning date of the range for messages to be displayed. All messages meeting the time range criteria that occurred after this date are displayed. The default is the current date. This parameter is optional.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To display information beginning with messages created a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE= -7.

Value	Description	Example
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for messages to be displayed. All messages meeting the time range criteria that occurred after this time are displayed. If you do not specify time, all messages that occurred in the last hour are displayed. You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with BEGINTIME=NOW+3 or BEGINTIME=+3, IBM Spectrum Protect™ displays messages with a time of 12:00 or later on the begin date.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-04:00 <i>or</i> -04:00. If you issue the QUERY ACTLOG command at 9:00 with BEGINTime=NOW-3:30 or BEGINTime= -3:30, IBM Spectrum Protect displays messages with a time of 5:30 or later on the begin date.

ENDDate

Specifies the ending date of the range for messages to be displayed. All messages meeting the time range criteria that occurred before this date are displayed. If you do not specify a value, the current date is used. This parameter is optional. You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 <i>or</i> -1. To display information created up to yesterday, you can specify ENDDATE=TODAY-1 or simply ENDDATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for messages to be displayed. All messages meeting this time range criteria that occurred before this time are displayed. If you do not specify a value, all messages are displayed up to the time when you issued this command. This parameter is optional.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, IBM Spectrum Protect displays messages with a time of 12:00 or earlier on the end date you specify.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 <i>or</i> -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, IBM Spectrum Protect displays messages with a time of 5:30 or earlier on the end date you specify.

MSGno

Specifies an integer that defines the number of the message to be displayed from the activity log. This integer is just the numeric part of the message. This parameter is optional.

Search

Specifies a text string that you want to search for in the activity log. Enclose the string expression in quotation marks if it contains blanks. You can use text and a wildcard character to specify this string. This parameter is optional.

Note: Do not enter as a text string either the IBM Spectrum Protect server name or text and a wildcard character that would find the server name. If you do so, the output includes messages that do not include the search string.

NODName

Specifies that the query displays messages logged for this node. If you do not specify a value for this parameter, messages for all nodes are displayed.

ORiginator

Specifies that the query displays messages logged by the server, client, or both. The default is ALL. Possible values are:

ALL

Specifies that the query displays messages that originated from the client and the server.

SErver

Specifies that the query displays messages that originated from the server.

CLient

Specifies that the query displays messages that originated from the client.

You can specify one of the following values to minimize processing time when querying the activity log for messages logged by the client:

OWNERname

Specifies that the query displays messages logged for a particular owner. If you do not specify a value for this parameter, messages for all owners are displayed.

SCHedname

Specifies that the query displays messages logged by a particular scheduled client activity. If you do not specify a value for this parameter, messages for all schedules are displayed.

DOmainname

Specifies that the query displays messages logged for a particular policy domain to which a named schedule belongs. This parameter is optional, unless you are specifying a schedule name.

SESSnum

Specifies that the query displays messages logged from a particular client session number. If you do not specify a value for this parameter, messages for all client sessions are displayed.

Example: Search activity log for messages with specific text

Search the activity log for any message that contains the string "delete". The output includes only messages produced during the past hour. Issue the command:

```
query actlog search=delete
```

Date/Time	Message
08/27/1998 15:19:43	ANR0812I Inventory client file expiration complete: 0 files deleted.

Example: Search activity log for messages within a specific time frame

Display messages that occurred yesterday between 9:30 and 12:30. Issue the command:

```
query actlog begindate=today-1  
begintime=09:30:00 endtime=12:30:00
```

Date/Time	Message
10/21/1998 10:52:36	ANR0407I Session 3921 started for administrator ADMIN (WebBrowser) (HTTP 9.115.20.100(2315)).
10/21/1998 11:06:08	ANR0405I Session 3922 ended for administrator ADMIN (WebBrowser).
10/21/1998 12:16:50	ANR0405I Session 3934 ended for administrator ADMIN (WebBrowser).

Example: Search activity log for messages from a specific client node

Search the activity log for IBM Spectrum Protect messages from the client for node JEE. Issue the command:

```
query actlog originator=client node=jee
```

Date/Time	Message
06/10/1998 15:46:22	ANE4007E (Session No: 3 Node: JEE) Error processing '/jee/report.out': access to the object is denied
06/11/1998 15:56:56	ANE4009E (Session No: 4 Node: JEE) Error processing '/jee/work.lst': disk full condition

Example: Search activity log for client and server messages from a specific client node and session

Search the activity log for IBM Spectrum Protect messages from the client and server for node A associated with Session 1. The output includes all messages with the defined text string, "SESSION: 1". Issue the command:

```
query actlog search="(SESSION:1)"
```

Date/Time	Message
02/13/2012 12:13:42	ANR0406I Session 1 started for node A (WinNT) (Tcp/Ip colind(2463)). (SESSION: 1)
02/13/2012 12:13:56	ANE4952I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects inspected: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4954I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects backed up: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4958I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects updated: 0 (SESSION: 1)
02/13/2012 12:13:56	ANE4964I (ANE4985I Session: 1, ANE4986I Node: A) Elapsed processing time: 00:00:02 (SESSION: 1)

02/13/2012 12:13:59 ANR0403I Session 1 ended for node A (WinNT).
(SESSION: 1)

Example: Search activity log for client-generated messages from a client session

Search the activity log for IBM Spectrum Protect messages from a specific client session. The output includes only messages generated by the client. Issue the command:

```
query actlog sessnum=1
```

```
Date/Time          Message
-----
02/13/2012 12:13:56 ANE4952I (ANE4985I Session: 1, ANE4986I Node: A)
                    Total number of objects inspected:      34
                    (SESSION: 1)
02/13/2012 12:13:56 ANE4954I (ANE4985I Session: 1, ANE4986I Node: A)
                    Total number of objects backed up:      34
                    (SESSION: 1)
02/13/2012 12:13:56 ANE4958I (ANE4985I Session: 1, ANE4986I Node: A)
                    Total number of objects updated:        0
                    (SESSION: 1)
02/13/2012 12:13:56 ANE4964I (ANE4985I Session: 1, ANE4986I Node: A)
                    Elapsed processing time:                00:00:02
                    (SESSION: 1)
```

Field descriptions

Date/Time

Specifies the date and time when the message was generated by the server or client.

Message

Specifies the message that was generated by the server or client.

Related commands

Table 1. Command related to QUERY ACTLOG

Command	Description
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.

QUERY ADMIN (Display administrator information)

Use this command to display information about one or more administrators.

Privilege class

Any administrator can issue this command.

Syntax

```
..*-----
>>-Query Admin--+----->
                    '-admin_name-'

>--+----->
|               .-,-----|
|               V         |
| '-Classes-----+System-----+'
|               +-Policy---+
|               +-Storage--+
|               +-Operator+
|               '-Node-----'

.-Format-----Standard-----
>--+----->
```

```
'-Format-----+Standard--+'
      '-Detailed-'
>--+-----+-----+-----+-----><
      '-AUTHentication---+LOcal--+'   '-ALerts-----+Yes--+'
              '-LDap--'                '-No--'
```

Parameters

admin_name

Specifies the name of the administrator for which you want to display information. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all administrators are displayed.

Classes

Specifies that you want to restrict output to those administrators that have privilege classes that you specify. This parameter is optional. You can specify multiple privilege classes in a list by separating the names with commas and no intervening spaces. If you do not specify a value for this parameter, information about all administrators is displayed, regardless of privilege class. Possible values are:

System

Display information on administrators with system privilege.

Policy

Display information on administrators with policy privilege.

Storage

Display information on administrators with storage privilege.

Operator

Display information on administrators with operator privilege.

Node

Display information on users with client node privilege.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified administrators.

Detailed

Specifies that complete information is displayed for the specified administrators.

Authentication

Specifies the password authentication method for the administrator.

Local

Display those administrators authenticating to the IBM Spectrum Protect™ server.

LDap

Display those administrators authenticating to an LDAP directory server. The administrator password is case-sensitive.

Alert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No

Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the QUERY MONITORSETTINGS command.

Example: Display information about all administrators

Display partial information on all administrators. Issue the command:

```
query admin
```

```
Administrator   Days Since   Days Since   Locked?   Privilege Classes
Name            Last Access  Password
```

		Set		
ADMIN	<1	<1	No	System
SERVER_CONSOLE			No	System

See Field descriptions for field descriptions.

Example: Display complete information about one administrator

From a managed server, display complete information for the administrator named ADMIN. Issue the command:

```
query admin admin format=detailed
```

```
Administrator Name: ADMIN
Last Access Date/Time: 1998.06.04 17.10.52
Days Since Last Access: <1
Password Set Date/Time: 1998.06.04 17.10.52
Days Since Password Set: 26
Invalid Sign-on Count: 0
Locked?: No
Contact:
System Privilege: Yes
Policy Privilege: **Included with system privilege**
Storage Privilege: **Included with system privilege**
Operator Privilege: **Included with system privilege**
Client Access Privilege: **Included with system privilege**
Client Owner Privilege: **Included with system privilege**
Registration Date/Time: 05/09/1998 23:54:20
Registering Administrator: SERVER_CONSOLE
Managing profile:
Password Expiration Period: 90 Day (s)
Email Address:
Email Aerts: Yes
Authentication: Local
SSL Required: No
Session Security: Strict
Transport Method: TLS 1.2
```

See Field descriptions for field descriptions.

Field descriptions

Administrator Name

Specifies the name of the administrator.

Last Access Date/Time

Specifies the date and time that the administrator last accessed the server.

Days Since Last Access

Specifies the number of days since the administrator last accessed the server.

Password Set Date/Time

Specifies the date and time that the administrator's password was defined or most recently updated.

Days Since Password Set

Specifies the number of days since the administrator's password was defined or most recently updated.

Invalid Sign-on Count

Specifies the number of invalid sign-on attempts that have been made since the last successful sign-on. This count can only be non-zero when an invalid password limit (SET INVALIDPWLIMIT) is greater than zero. When the number of invalid attempts equals the limit set by the SET INVALIDPWLIMIT command, the administrator is locked out of the system.

Locked?

Specifies whether the administrator is locked out of the system.

Contact

Specifies any contact information for the administrator.

System Privilege

Specifies whether the administrator has been granted system privilege.

Policy Privilege

Specifies whether the administrator has been granted unrestricted policy privilege or the names of any policy domains that the restricted policy administrator can manage.

Storage Privilege

Specifies whether the administrator has been granted unrestricted storage privilege or the names of any storage pools that the restricted storage administrator can manage.

Operator Privilege

Specifies whether the administrator has been granted operator privilege.

Client Access Privilege

Specifies that client access authority has been granted to a user with node privilege.

Client Owner Privilege

Specifies that client owner authority has been granted to a user with node privilege.

Registration Date/Time

Specifies the date and time that the administrator was registered.

Registering Administrator

Specifies the name of the administrator who registered the administrator. If this field contains `$$CONFIG_MANAGER$$`, the administrator is associated with a profile that is managed by the configuration manager.

Managing Profile

Specifies the profiles to which the managed server subscribed to get the definition of this administrator.

Password Expiration Period

Specifies the administrator's password expiration period.

Email Address

Specifies the email address for the administrator.

Email Alerts

Specifies whether alerts are sent to the specified administrator by email.

Authentication

Specifies the password authentication method: LOCAL, LDAP, or LDAP (pending).

Authentication Target	Authentication Method
IBM Spectrum Protect server	LOCAL
LDAP directory server	LDAP
This administrator is configured to authenticate with an LDAP directory server, but the administrator did not yet authenticate through a client node.	LDAP (pending)

SSL Required (deprecated)

Specifies whether the security setting for the administrator user ID requires the Secure Sockets Layer (SSL) protocol. Values can be YES, NO, or Default. You must have system level authority to update the administrator SSLREQUIRED setting. This parameter is deprecated.

Session Security

Specifies the level of session security that is enforced for the administrator ID. Values can be STRICT or TRANSITIONAL.

Transport Method

Specifies the transport method that was last used for the specified administrator. Values can be TLS 1.2, TLS 1.1, or NONE. A question mark (?) is displayed until a successful authentication is completed.

Related commands

Table 1. Commands related to QUERY ADMIN

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REMOVE ADMIN	Removes an administrator from the list of registered administrators.
RENAME ADMIN	Changes an IBM Spectrum Protect administrator's name.
RESET PASSEXP	Resets the password expiration for nodes or administrators.

Command	Description
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.
SET MINPWLENGTH	Sets the minimum length for client passwords.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.

QUERY ALERTTRIGGER (Query the list of defined alert triggers)

Use this command to display which server messages are defined as alerts.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query ALERTTrigger-----*----->>
      |-----+-----|
      |---message_number---|
```

Parameters

message_number

Specifies the message number that you want to query. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers. If you do not specify a message number, all alert triggers are displayed.

Query alert triggers to display which messages are designated as alerts

Display all messages that are designated as alerts by issuing the following command:

```
query alerttrigger
```

Example output:

Alert Trigger	Category	Administrator
ANR1067E	SERVER	HARRYH
ANR1073E	SERVER	CSDADMIN, DJADMIN, HARRYH
ANR1074E	STORAGE	CSDADMIN, DJADMIN, HARRYH
ANR1096E	STORAGE	CSDADMIN, DJADMIN, HARRYH, MHAYE

Query alert triggers for a specific message number

Display all alert triggers that have message number ANR1067E designated to them by issuing the following command:

```
query alerttrigger ANR1067E
```

Example output:

Alert Trigger	Category	Administrator
ANR1067E	SERVER	HARRYH

Field descriptions

Alert Trigger

Status

Specifies the status type that you want to display. If you do not specify a status, all alerts are queried and displayed. Specify one of the following values:

Active

Displays alerts that are specified in the IBM Spectrum Protect server database as active.

INactive

Displays alerts that are in the inactive state.

CLosed

Displays alerts that are in the closed state.

ANy

Displays all alerts, without regard to state.

MSGnum

Specifies the message number that you want to display. Specify the numerical portion of an IBM Spectrum Protect server message. Values are in the range 0 - 9999. For example, the message number in message ANR2044E is 2044. Specify multiple message numbers by separating them with commas and no intervening spaces.

CATegory

Specifies the category type for the alert, which is determined by the message types. Specify one of the following values:

APplication

Alert is classified as application category. For example, you can specify this category for messages that are associated with application (TDP) clients.

INventory

Alert is classified as inventory category. For example, you can specify this category for messages that are associated with the database, active log file, or archive log file.

Note: The category of `CATalog` is used instead of `INventory` in alerts from servers that were not upgraded to IBM Spectrum Protect 7.1.0 or later.

CLient

Alert is classified as client category. For example, you can specify this category for messages that are associated with general client activities.

DEvice

Alert is classified as device category. For example, you can specify this category for messages that are associated with device classes, libraries, drives, or paths.

SErver

Alert is classified as general server category. For example, you can specify this category for messages that are associated with general server activities or events.

STorage

Alert is classified as storage category. For example, you can specify this category for messages that are associated with storage pools.

SYstems

Alert is classified under system clients category. For example, you can specify this category for messages that are associated with system backup and archive or hierarchical storage management (HSM) backup-archive clients.

VMclient

Alert is classified under VMclient category. For example, you can specify this category for messages that are associated with virtual machine clients.

SOURCEType

Specifies the source type that is being queried. Specify one of the following values:

LOcal

Displays alerts that originated from the local IBM Spectrum Protect server.

CLient

Displays alerts that originated from the IBM Spectrum Protect client.

REmote

Displays alerts that originated from another IBM Spectrum Protect server.

SOURCENAME

Specifies the name of the source where the alert originated. SOURCE_NAME can be the name of a local or remote IBM Spectrum Protect server, or an IBM Spectrum Protect client.

ID

This optional parameter specifies the unique ID of the alert that you want to display. Specify a value from 1 to 9223372036854775807.

ASSigned

Specifies the administrator name that is assigned the alert that you want to query.

RESolvedby

Specifies the administrator name that resolved the alert that you want to query.

Query active alerts

Display only alerts that are active in the server database by issuing the following command:

```
query alertstatus status=active
```

Query active alerts for two messages issued by the local server

Issue the following command to display only active alerts for message numbers ANE4958I and ANR4952E that were issued by the local server:

```
query alertstatus msgnum=4958,4952 status=active sourcetype=local
```

Query active alerts for messages ANR4958I and ANR4952E issued by a client

Issue the following command to display only active alerts for message numbers ANE4958I and ANE4952I that were issued by a client:

```
query alertstatus msgnum=4958,4952 status=active sourcetype=client
```

Query all alerts on a server

Issue the following command to display all alerts that are on the server:

```
query alertstatus
```

Example output: Display all the alerts that are on the server:

```
Alert Identifier: 83
Alert Message Number: 293
  Source Name: SEDONA
  Source Type: LOCAL
  First Occurrence: 03/07/2013 17:08:35
Most Recent Occurrence: 03/07/2013 17:08:35
  Count: 1
  Status: ACTIVE
  Last Status Change: 12/31/1969 17:00:00
  Category: INVENTORY
  Message: ANR0293I Reorganization for table AF_BITFILES
          started.
  Assigned:
  Resolved By:
  Remark:
```

```
Alert Identifier: 85
Alert Message Number: 293
  Source Name: SEDONA
  Source Type: LOCAL
  First Occurrence: 03/08/2013 05:45:00
Most Recent Occurrence: 03/08/2013 05:45:00
  Count: 1
  Status: ACTIVE
  Last Status Change: 12/31/1969 17:00:00
  Category: INVENTORY
  Message: ANR0293I Reorganization for table
          BF_AGGREGATED_BITFILES started.
  Assigned:
  Resolved By:
  Remark:
```

```
Alert Identifier: 1282
Alert Message Number: 293
  Source Name: ALPINE
  Source Type: LOCAL
  First Occurrence: 02/13/2013 15:47:50
```

Most Recent Occurrence: 02/13/2013 15:47:50
 Count: 1
 Status: CLOSED
 Last Status Change: 02/26/2013 09:46:39
 Category: INVENTORY
 Message: ANR0293I Reorganization for table
 TSMON_ALERT started.
 Assigned:
 Resolved By:
 Remark:

Alert Identifier: 1792
 Alert Message Number: 293
 Source Name: ALPINE
 Source Type: LOCAL
 First Occurrence: 02/19/2013 08:58:14
 Most Recent Occurrence: 02/19/2013 08:58:14
 Count: 1
 Status: CLOSED
 Last Status Change: 03/01/2013 12:39:21
 Category: INVENTORY
 Message: ANR0293I Reorganization for table
 ACTIVITY_LOG started.
 Assigned:
 Resolved By:
 Remark:

Field descriptions

Alert Identifier

The unique identifier for the alert.

Alert Message Number

The message number for the alert.

Source Name

The name of the source from where the alert originated.

Source Type

The type of the originating source.

First Occurrence

The date and time when the alert first occurred.

Most Recent Occurrence

The date and time when the alert occurred last.

Count

The total number of times the alert has been triggered.

Status

Specifies the status of the alert.

Last Status Change

Specifies the time and date when the status for the alert last changed.

Category

The category for the alert.

Message

The message that triggers the alert.

Assigned

Specifies the user whom this alert concerns.

Resolved By

Species the user who has investigated and resolved the alert.

Remark

An optional remark to be left by the resolver.

Related commands

Table 1. Commands related to QUERY ALERTSTATUS

Command	Description
DEFINE ALERTTRIGGER (Define an alert trigger)	Associates specified messages to an alert trigger.

Related commands

Table 1. Commands related to QUERY ASSOCIATION

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DELETE ASSOCIATION	Deletes the association between clients and a schedule.

QUERY AUDIT OCCUPANCY (Query client node storage utilization)

Use this command to display information about client node server storage utilization. To display current license audit information from the server, use the AUDIT LICENSE command before you issue the QUERY AUDIT OCCUPANCY command.

As part of a license audit operation, the server calculates, by node, the amount of backup, archive, and space management storage in use. For servers that manage large amounts of data, this calculation can take a great deal of processor time and can stall other server activity. You can use the AUDIT STORAGE server option to specify that storage is not to be calculated as part of a license audit.

You can use the information from this query to determine if and where client node storage utilization must be balanced. This information can also assist you with billing clients for storage usage.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query AUDITOccupancy----->
      | .-,----- . |
      | v          | |
      |---node_name---|
>----->
      | .-,----- . |
      | v          | |
      |---Domain-----domain_name---|
      .-Pooltype-----ANY-----
>----->>
      |---Pooltype-----ANY-----|
      +-Primary-+
      |---Copy-----|
```

Parameters

node_name

Specifies a list of nodes for which to display server storage use information. Specify more than one node by separating the node names with commas, with no intervening spaces. You can use wildcard characters to specify names. The default (*) is to query all client nodes. Use the DOMAIN parameter to limit this list by policy domain. This parameter is optional.

Domain

Specifies a list of policy domains to restrict which nodes are displayed. Nodes belonging to the specified policy domains are displayed. Specify more than one policy domain by separating the policy domain names with commas, with no intervening spaces. You can use wildcard characters to specify names. This parameter is optional.

Pooltype

Specifies the type of storage pool to display. This parameter is optional. The default is ANY. Possible values are:

ANY

Specifies both primary and copy storage pools. The value that is presented is the total for the two pools.

Primary

Specifies primary storage pools only.

COPY

Specifies copy storage pools only.

Example: Display storage usage

Display combined storage use in primary and copy storage pools. Issue the command:

```
query auditoccupancy
```

License information as of last audit on 05/22/1996 14:49:51.

Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
CLIENT	245	20	0	265
SMITH	245	20	0	265
SMITHERS	245	20	0	265
JOHNSON	300	15	0	320
JOE	245	20	0	265
TOM	300	15	0	320
LARRY	245	20	0	265

See Field descriptions for field descriptions.

Field descriptions

Node Name

Specifies the name of the client node.

Backup Storage Used (MB)

Specifies the total backup storage use for the node. For this value, one MB = 1048576 bytes.

Archive Storage Used (MB)

Specifies the total archive storage use for the node. For this value, one MB = 1048576 bytes.

Space-Managed Storage Used (MB)

Specifies the amount of server storage that is used to store files that are migrated from the client node by an IBM Spectrum Protect™ for Space Management client. For this value, one MB = 1048576 bytes.

Total Storage Used (MB)

Specifies the total storage use for the node. For this value, one MB = 1048576 bytes.

Related commands

Table 1. Commands related to QUERY AUDITOCCUPANCY

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY LICENSE	Displays information about licenses and audits.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the IBM Spectrum Protect server.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.

QUERY BACKUPSET (Query a backup set)

Use this command to display information about one or more backup sets.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query BACKUPSET----->
      .-*-----
      | .-,-----
      | V          |
      |-----node_name-----|
      |-----node_group_name-|

      .-*-----
>+----->
      | .-,-----
      | V          |
      |-----backup_set_name-|

>+----->
      '-BEGINTime----time-' '-ENDDate----date-'

>+----->
      '-ENDTime----time-' '-WHERERetention----days----'
                          '-NOLimit-'

>+----->
      '-WHEREDESCRIPTION----description-'

>+----->
      '-WHEREDEVclass----device_class_name-'

>+----->
      '-WHERETOExists----Yes--'
                          '-No--'

>+----->
      | .-,-----
      | V          |
      |-----WHEREDATAType-----FILE--|
      |-----IMAGE-|

      .-Format----Standard-----
>+-----><
      '-Format----Standard--'
      '-Detailed-'

```

Parameters

node_name or node_group_name

Specifies the name of the client node and node groups whose data is contained in the backup set to be displayed. To specify multiple node names and node group names, separate the names with commas and no intervening spaces. You can use wildcard characters with node names but not with node group names.

backup_set_name

Specifies the name of the backup set whose information is to be displayed. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

BEGINDate

Specifies the beginning date of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the BEGINTIME parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

Value	Description	Example
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the BEGINDATE parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes specified	NOW+02:00 <i>or</i> +02:00.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes specified	NOW-02:00 <i>or</i> -02:00.

ENDDate

Specifies the ending date of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the ENDTIME parameter to specify an ending date and time. If you specify an end date without an end time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days <i>or</i> +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 <i>or</i> +3.
TODAY-days <i>or</i> -days	The current date minus days specified.	TODAY -3 <i>or</i> -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the ENDDATE parameter to specify a date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
-------	-------------	---------

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes specified	NOW+02:00 <i>or</i> +02:00.
NOW-HH:MM <i>or</i> - HH:MM	The current time minus hours and minutes specified	NOW-02:00 <i>or</i> -02:00.

WHERERetention

Specifies the retention value, specified in days, that must be associated with the backup sets to be displayed. You can specify an integer from 0 to 30000. The values are:

days

Specifies that backup sets that are retained this number of days are displayed.

NOLimit

Specifies that backup sets that are retained indefinitely are displayed.

WHEREDescription

Specifies the description that must be associated with the backup set to be displayed. The description you specify can contain wildcard characters. This parameter is optional. Enclose the description in quotation marks if it contains any blank characters.

WHEREDEVclass

Specifies the name of the device class that must be associated with the backup set to be displayed. You can use wildcard characters to specify a device class name. This parameter is optional.

WHERETOExists

Specifies whether a backup set must have a table of contents in order to be displayed. This parameter is optional. The default is to display all backup sets whether or not they have a table of contents.

WHEREDATType

Specifies the data type of a backup set to be displayed. This parameter is optional. The default is to display all types of backup sets. To specify multiple data types, separate data types with commas and no intervening spaces.

FILE

Specifies that a file level backup set is to be displayed. File level backup sets contain files and directories backed up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be displayed. Image backup sets contain images created by the backup-archive client BACKUP IMAGE command.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified backup sets.

Detailed

Specifies that complete information is displayed for the specified backup sets.

Example: Query a backup set

Display information for backup sets whose names begin with PERS_DATA. The backup sets belong to the node JANE and are assigned to the DVLMENT device class.

```
query backupset jane pers_data*
      Node Name: JANE
      Backup Set Name: PERS_DATA.3089
      Data Type: File
      Date/Time: 03/17/2007 16:17:47
      Retention Period: 60
      Device Class Name: DVLMENT
      Description: backupset created from /srvr
Has Table of Contents (TOC)?: Yes
```

Field descriptions

- Node Name**
Specifies the name of the client node whose data is contained in the backup set.
- Backup Set Name**
Specifies the name of the backup set.
- Data Type**
Displays the data type of the backup sets. Possible types are file, image, and application.
- Date/Time**
Specifies the date and time (PITDate and PITTime) of the GENERATE BACKUPSET command. The PITDate and PITTime specify that files that were active on the specified date and time and that are still stored on the IBM Spectrum Protect™ server are to be included in the backup set, even if they are inactive at the time you issue the GENERATE BACKUPSET command. The default is the date on which the GENERATE BACKUPSET command is run.
- Retention Period**
Specifies the number of days that the backup set is retained on the server.
- Device Class Name**
Specifies the name of the device class for which the volumes containing the backup set is assigned.
- Description**
Specifies the description associated with the backup set.
- Has Table of Contents (TOC)?**
Specifies whether the backup set has a table of contents.

Related commands

Table 1. Commands related to QUERY BACKUPSET

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

QUERY BACKUPSETCONTENTS (Query contents of a backup set)

Use this command to display information about the files and directories contained in a backup set for a client node.

Remember: Processing this command can use considerable network resources and mount points.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```
>>-Query BACKUPSETCONTENTS--node_name--backup_set_name----->
    .-DATAType--===FILE----- .
>--+-----+----->>
    '-DATAType--===+FILE---+'
```

Parameters

node_name (Required)

Specifies the name of the client node whose data is contained in the backup set to display. The name you specify cannot contain wildcard characters nor can it be a list of node names separated by commas.

backup_set_name (Required)

Specifies the name of the backup set to display. The name that you specify cannot contain wildcard characters nor can it be a list of node names that are separated by commas.

DATATYPE

Specifies that the backup set containing the specified types of data is to be queried. This parameter is optional. The default is that a file level backup set is to be queried. Possible values are:

FILE

Specifies that a file level backup set is to be queried. File level backup sets contain files and directories backed up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be queried. Image backup sets contain images created by the backup-archive client BACKUP IMAGE command.

Example: Query contents of a backup set for a specific node

Display the contents from backup set named PERS_DATA.3099 belonging to client node JANE. Issue the command:

```
query backupsetcontents jane pers_data.3099
```

Node Name	Filespace Name	Client's Name for File
JANE	/svr	/deblock
JANE	/svr	/deblock.c
JANE	/svr	/dsmerror.log
JANE	/svr	/dsmxxxxx.log
JANE

Field descriptions

Node Name

Specifies the name of the client node whose data is contained in the backup set.

Filespace Name

Specifies the name of the file space to which the specified file belongs.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Client's Name for File

Specifies the name of the file.

File space names and file names that can be in a different code page or locale than the server do not display correctly in the Operations Center or the administrative command-line interface. The data itself is backed up and can be restored properly, but the file space or file name may display with a combination of invalid characters or blank spaces.

If the file space name is Unicode enabled, the name is converted to the server's code page for display. The results of the conversion for characters not supported by the current code page depends on the operating system. For names that IBM Spectrum Protect™ is able to partially convert, you may see question marks (??), blanks, unprintable characters, or "...".

These characters indicate to the administrator that files do exist. If the conversion is not successful, the name is displayed as "...". Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

A file name that is displayed as "....." indicates that both the file path and file name were not successfully converted. An example of the path and name could be:

```
my\dir\...
```

Related commands

Table 1. Commands related to QUERY BACKUPSETCONTENTS

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
DELETE BACKUPSET	Deletes a backup set.
QUERY BACKUPSET	Displays backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.

AIX Linux Windows

QUERY CLEANUP (Query the cleanup that is required in a source storage pool)

Use this command to display information about damaged files that are identified during a storage pool conversion process.

When you issue the CONVERT STGPOOL command to convert a FILE device class, a tape device class, or a virtual tape library (VTL) to a directory-container storage pool, some files in the source storage pool might not convert because of damaged data. To display damaged data that is identified during the conversion process, issue the QUERY CLEANUP command on a source storage pool.

To recover an undamaged version of the data from a copy or active-data storage pool, issue the RESTORE STGPOOL command. To recover an undamaged version of the data from a target replication server issue the REPLICATE NODE command and specify the RECOVERDAMAGED=YES parameter.

Privilege class

To issue this command, you must have restricted storage privilege.

Syntax

```
>>-Query CLeaup--pool_name-----><
```

Parameters

pool_name(Required)
Specifies the storage pool to query.

Example: Display damaged files that are identified by a storage pool conversion process

Display damaged files in a storage pool that is named POOL1. See Field descriptions for field descriptions.

```
query cleanup pool1  
  
File Name: \RTC\BDAT\GIGFILES\BF1.GB  
State: Active  
Stored Size: 1 GB
```

Filespace Name: \\ibm838-r90gf0gx\c\$
Type: Backup
Client Name: CAKINProtection
Protection Date: 03/25/2016 16:47:57

Field descriptions

File Name

The name of the damaged file.

State

The state of the data in the inventory. The following states are possible:

Active

The version of the file in the inventory is active. You can have only one active version of the file in the inventory.

Inactive

The version of the file in the inventory is inactive. You can have multiple inactive versions of the file in the inventory.

Stored Size

The size of the data, in megabytes (MB) or gigabytes (GB), that is stored in the storage pool.

Filespace Name

The name of the file space where the file is assigned.

Type

The type of operation that was used to store the file. The following types are possible:

Backup

Files that are backed up.

Archive

Files that are archived.

SpaceMg

Files that are migrated from an IBM Spectrum Protect™ for Space Management client.

Client Name

The name of the client that owns the file.

Protection Date

The time and date that the file was backed up, archived, or migrated by an IBM Spectrum Protect for Space Management client.

Related commands

Table 1. Commands related to QUERY CLEANUP

Command	Description
CONVERT STGPOOL	Convert a storage pool to a directory-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
QUERY CONVERSION	Query conversion status of a storage pool.
REMOVE DAMAGED	Removes damaged data from a source storage pool.
REPAIR STGPOOL	Repairs a directory-container storage pool.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.

QUERY CLOPTSET (Query a client option set)

Use this command to query a client option set.

Privilege class

Any administrator can issue this command.

Syntax

```
      .-*-----.  
>>-Query CLOptset-----+-----+----->  
      '-option_set_name-'  
  
>--+-----+----->>  
      '-DEscription----description-'
```

Parameters

option_set_name

Specifies the name of the client option set to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is option set names.

DEscription

Specifies the description used on the DEFINE or UPDATE CLOPTSET commands to be used as a filter. If the description contains spaces, enclose it in quotation marks. This parameter is optional.

Example: Query a client option set

From a managed server, query a client option set named ENG. Issue the following command:

```
query cloptset eng  
  
      Optionset:  ENG  
      Description:  
Last Update by (administrator): $$CONFIG_MANAGER$$  
      Managing profile:  
      Replica Option Set: Yes  
  
      Option: SCROLLINES  
      Sequence number: 0  
Use Option Set Value (FORCE): No  
      Option Value: 40  
  
      Option: SCROLLPROMPT  
      Sequence number: 0  
Use Option Set Value (FORCE): No  
      Option Value: yes
```

Field descriptions

Optionset

Specifies the name of the option set.

Description

Specifies the description of the client option set.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the option set. If this field contains \$\$CONFIG_MANAGER\$\$, the client option set is associated with a profile that is managed by the configuration manager.

Managing profile

Specifies the profile to which the managed server subscribed to get the definition of the client option set.

Replica Option Set

Specifies the replica option set is replicated by the source replication server.

Option

Specifies the name of the option.

Sequence number

Specifies the sequence number of the option.

Use Option Set Value (FORCE)

Specifies whether the server option setting overrides the option setting for the client. NO indicates that the server option setting does not override the client option. YES indicates that the server option setting overrides the client option setting. This option is set with the FORCE parameter on the DEFINE CLIENTOPT command.

Option Value

Specifies the value of the option.

Related commands

Table 1. Commands related to QUERY CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.
DEFINE PROFASSOCIATION	Associates objects with a profile.

QUERY COLLOGROUP (Query a collocation group)

Use this command to display the collocation groups defined on the server.

Privilege class

Any administrator can issue this command.

Syntax

```
.-*-----.  
>>-Query COLLOGGroup--+-----+----->  
      '-group_name-'  
  
.-Format----Standard----.  
>--+-----+----->>  
      '-Format----+Standard-+-'  
      '-Detailed-'
```

Parameters

group_name

Specifies the name of the collocation group to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all collocation groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed. To display the members of the collocation group, you must specify FORMAT=DETAILED.

Display defined collocation groups

Display the collocation groups defined on the server. Issue the following command:

```
query collogroup
```

```
Collocation Group Name      Collocation Group Description  
-----
```

DEPT_ED Education department
GROUP1 Low cap client nodes.

See Field descriptions for field descriptions.

Display detailed information for collocation groups

Display complete information about all collocation groups and determine which client nodes belong to which collocation groups. Issue the following command:

```
query collogcgroup format=detailed

    Collocation Group Name: DEPT_ED
    Collocation Group Description: Education department
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2013 10:59:03
    Collocation Group Member(s): EDU_1 EDU_7
    Filespace Member(s):

    Collocation Group Name: GROUP1
    Collocation Group Description: Low cap client nodes.
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2013 10:59:16
    Collocation Group Member(s): CHESTER
    Filespace Member(s): alpha

    Collocation Group Name: GROUP1
    Collocation Group Description: Low cap client nodes.
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2013 10:59:16
    Collocation Group Member(s): CHESTER
    Filespace Member(s): beta

    Collocation Group Name: GROUP1
    Collocation Group Description: Low cap client nodes.
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2013 10:59:16
    Collocation Group Member(s): CHESTER
    Filespace Member(s): gamma
```

See Field descriptions for field descriptions.

Field descriptions

Collocation Group Name

The name of the collocation group.

Collocation Group Description

The description for the collocation group.

Last Update by (administrator)

The name of the administrator that defined or most recently updated the collocation group.

Last Update Date/Time

The date and time that an administrator defined or most recently updated the collocation group.

Collocation Group Member(s)

The members of the collocation group.

Filespace Member(s)

The file space or file spaces that are members of the collocation group. If there is more than one file space, each file space is displayed in a separate entry.

Related commands

Table 1. Commands related to QUERY COLLOGCROUP

Command	Description
DEFINE COLLOGCROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.

STGpool

Specifies the name of the directory-container storage pool. This parameter is optional. The maximum length of the storage pool name is 30.

Format

Specifies the level of detail of the query results. This parameter is optional. Specify one of the following values:

Standard

Specifies that a summary of the information is displayed. This value is the default.

Detailed

Specifies that detailed information is displayed.

State

Specifies the state of the container that is queried. This parameter is optional. Specify one of the following values:

AVAILABLE

Specifies that only containers that are available are displayed.

UNAVAILABLE

Specifies that only containers that are not available are displayed. For example, a container might be unavailable if the header is corrupted or if the container cannot be opened.

ANY

Specifies that containers in any state are displayed. This value is the default.

READONLY

Specifies that only containers in a read-only state are displayed. Data in the container can be read but data cannot be written to the container.

PENDING

Specifies that only containers in a pending state are displayed.

TYPE

Specifies the type of container that is queried. This parameter is optional. Specify one of the following values:

NONDEDUP

Displays containers that contain data that is not deduplicated. This type of data includes metadata, encrypted data, and data that is too small for data deduplication.

DEDUP

Displays containers that contain deduplicated data.

CLOUD

Displays containers that are stored in a cloud storage pool.

ANY

Displays any type of container. This value is the default.

AIX | Linux

Example: Display information about a container

See Field descriptions for field descriptions.

```
query container /Containers/09/0000000000000943.ncf
```

Container	Storage Pool Name	Container Type	State
/Containers/09/0000000000000943.ncf	STGPOOL1	Non Dedup	Available

Windows

Example: Display information about a container

See Field descriptions for field descriptions.

```
query container C:\abc\00\0000000000000005.ncf
```

Container	Storage Pool Name	Container Type	State
C:\abc\00\0000000000000005.ncf	STGPOOL1	Non Dedup	Available

AIX | Linux

Example: Display detailed information about a container

Display detailed information about containers that contain deduplicated data in storage pool STGPOOL1:

```
query container stgpool=STGPOOL1 type=dedup format=detail

        Container: /abc/00/0000000000000001.dcf
Storage Pool Name: STGPOOL1
  Container Type: Dedup
    State: Available
Maximum size (MB): 40,960
  Free Space (MB): 39,700
Approx. Date Last Written: 11/10/2014 15:17:09
Approx. Date Last Audit:
  Cloud Type:
    Cloud URL:
Cloud Object Size (MB):
Space Utilized (MB):
Data Extent Count:
```

Windows

Example: Display detailed information about a container

Display detailed information about containers that contain deduplicated data in storage pool STGPOOL1:

```
query container stgpool=STGPOOL1 type=dedup format=detail

        Container: C:\abc\00\0000000000000001.dcf
Storage Pool Name: STGPOOL1
  Container Type: Dedup
    State: Available
Maximum size (MB): 40,960
  Free Space (MB): 39,700
Approx. Date Last Written: 11/10/2014 15:17:09
Approx. Date Last Audit:
  Cloud Type:
    Cloud URL:
Cloud Object Size (MB):
Space Utilized (MB):
Data Extent Count:
```

Example: Display detailed information about containers that are stored in a cloud storage pool

Display detailed information about containers that are stored in the cloud storage pool CLOUDPOOL:

```
query container stgpool=CLOUDPOOL format=detail

        Container: 7-64a1261000c811e58e8f005056c00008
Storage Pool Name: CLOUDPOOL
  Container Type: Cloud
    State:
Free Space (MB):
Maximum Size (MB):
Approx. Date Last Written: 05/22/2015 14:36:57
Approx. Date Last Audit:
  Cloud Type: SWIFT
    Cloud URL: http://cloudurl:5000/v2.0
Cloud Object Size (MB):
Space Utilized (MB): 27
Data Extent Count: 95
```

Field descriptions

Container

The name of the container.

Storage Pool Name

The name of the storage pool.

Container Type

The type of container.

State

The state of the data in the container. The field can contain one of the following values:

Available

The container is available for use.

Unavailable

The container cannot be opened or validated.

Tip: Issue the AUDIT CONTAINER command to validate the contents of the container.

Read only

The container can be read but data cannot be written to the container.

Pending

The container is pending deletion. When the value that is specified for the REUSEDELAY parameter expires on the DEFINE STGPOOL or UPDATE STGPOOL command, the container is deleted.

In general, this field does not apply to containers that are stored in cloud-container storage pools. However, if a container in a cloud-container storage pool is moved by using the MOVE CONTAINER command with the DEFRAG=YES setting, the container is in pending state until it is deleted.

Maximum Size (MB)

The maximum size of the container, in megabytes.

This field does not apply to containers that are stored in cloud storage pools.

Free Space (MB)

The total amount of free space that is available in the container, in megabytes.

This field does not apply to containers that are stored in cloud storage pools.

Approx. Date Last Written

The approximate date and time that data was written to the container.

Approx. Date Last Audit

The approximate date and time that data was audited in the container.

Cloud Type

If the container is stored in a cloud storage pool, the type of cloud platform.

Cloud URL

If the container is stored in a cloud storage pool, the URL for accessing the on-premises private cloud or off-premises public cloud.

Cloud Object Size (MB)

The size of the cloud object, in megabytes, if the container is represented by a single object in the cloud-container storage pool.

Space Utilized (MB)

If the container is stored in a cloud storage pool, the amount of space that is used by the container in the on-premises private cloud or off-premises public cloud.

Data Extent Count

If the container is stored in a cloud-container storage pool, the number of data extents that are managed by the on-premises private cloud or off-premises public cloud for the container.

Table 1. Commands related to QUERY CONTAINER

Command	Description
AUDIT CONTAINER	Audit a directory-container storage pool.
MOVE CONTAINER	Moves the contents of a storage pool container to another container.
QUERY DAMAGED	Displays information about damaged files.

QUERY CONTENT (Query the contents of a storage pool volume)

Use this command to display information about files in a storage pool volume, and the names of client files that link to a deduplicated group of files.

You can use this command to identify files that the server found to be damaged and files that were backed up to a copy storage pool or copied to an active-data pool. This command is useful when a volume is damaged or before you:

- Request the server to fix inconsistencies between a volume and the database
- Move files from one volume to another volume
- Delete a volume from a storage pool

Because this command can take a long time to run and the results can be large, consider using the COUNT parameter to limit the number of files displayed.

Note: Files that are cached in a disk volume and that are marked as damaged are not included in the results.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query CONtEnt--volume_name--+-----+----->
                                '-NODE---node_name-'
>--+-----+-----+----->
  '-Filespace---file_space_name-' '-COUnT---number-'

.-Type---ANY-----.-Format---Standard-----
>--+-----+-----+----->
  '-Type---+ANY-----+' '-Format---+Standard-+-'
      +-Backup-----+           '-Detailed-'
      +-Archive-----+
      '-Spacemanaged-'

                                (1)
.-DAmaged---ANY-----.-COPIed---ANY-.
>--+-----+-----+----->
  '-DAmaged---+ANY-+-' '-COPIed---+ANY-+-'
      +-Yes-+           +-Yes-+
      '-No--'           '-No--'

.-NAMEType---SERVER-----
>--+-----+-----+----->
  '-NAMEType---+SERVER-+-'
      +-UNICODE-+
      '-FSID----'

.-CODEType---BOTH-----
>--+-----+-----+----->
  '-CODEType---+UNICODE-----+'
      +-NONUNICODE-+
      '-BOTH-----'

.-FOLLOWLinks---No-----
>--+-----+-----+-----><
  '-FOLLOWLinks---+No-----+-'
      +-Yes-----+
      '-JUSTLinks-'
```

Notes:

1. Use this parameter only for volumes in primary storage pools.

Parameters

volume_name (Required)

Specifies the volume to be queried.

NODE

Specifies the backup-archive client or the IBM Spectrum Protect™ for Space Management associated with the file space to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a name, all backup-archive and IBM Spectrum Protect for Space Management clients are included.

Filespace

Specifies the file space to query. This parameter is optional. You can use wildcard characters to specify this name. File space names are case-sensitive. If you do not specify a file space name, all file spaces are included.

For a server that has clients with Unicode support, you might need to have the server convert the file space name that you enter. For example, you might need to have the server convert the name that you enter from the server's code page to Unicode. See the NAMETYPE parameter for details. If you do not specify a file space name or specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or non-Unicode file spaces.

COUnt

Specifies the number of files to be displayed. This parameter is optional. You can specify either a positive integer or a negative integer. If you specify a positive integer, *n*, the first *n* files are displayed. If you specify a negative integer, *-n*, the last *n* files are displayed in *reverse* order. You cannot specify COUNT=0. If you do not specify a value for this parameter, all files are displayed.

Type

Specifies the types of files to query. This parameter is optional. The default value is ANY. If the volume that is being queried is assigned to an active-data pool, the only valid values are ANY and BACKUP. Possible values are:

ANY

Specifies that all types of files in the storage pool volume are queried; backup versions of files, archived copies of files, and files that are migrated by IBM Spectrum Protect for Space Management clients from client nodes.

Backup

Specifies that only backup files are queried.

Archive

Specifies that only archive files are queried. This value is not valid for active-data pools.

SPacemanaged

Specifies that only space-managed files (files that were migrated by an IBM Spectrum Protect for Space Management client) are queried. This value is not valid for active-data pools.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed. Unicode names are converted to the server code page.

Detailed

Specifies that complete information is displayed. Unicode names are displayed in hexadecimal.

DAMaged

Specifies criteria to restrict the query output based on whether files are marked as damaged. For purposes of this criteria, the server examines only physical files (a file that might be a single logical file or an aggregate that consists of logical files). This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the server found the files to be damaged.

Yes

Specifies that only files that are marked as damaged are displayed. These are files in which the server found errors when a user attempted to restore, retrieve, or recall the file, or when an AUDIT VOLUME command was run.

No

Specifies that only files not known to be damaged are displayed.

COPIed

Specifies criteria to restrict the query output based on whether files were backed up to a copy storage pool. Whether files are stored in an active-data pool does not affect the output. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the files are backed up to a copy storage pool. Primary and cached file copies are displayed.

Yes

Specifies that the files displayed are only those for which at least one usable backup copy exists in a copy storage pool. A file is not displayed if its copy in the copy storage pool is known to have errors. Cached file copies are not displayed because these files are never restored.

Use COPIED=YES to identify primary files that can be restored using the RESTORE VOLUME or RESTORE STGPOOL command.

No

Specifies that the files displayed are only those for which no usable backup copies exist in a copy storage pool. Cached file copies are not displayed because these files are never restored.

Use COPIED=NO to identify primary files that cannot be restored using the RESTORE VOLUME or RESTORE STGPOOL command.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with Unicode support. A backup-archive client with Unicode support is currently available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare. Use this parameter only when you specify a partly or fully qualified file space name.

The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the names and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specify how you want the server to interpret the file space names that you enter. Use this parameter only when you enter a single wildcard character for the file space name.

The default value is BOTH, which means that the file spaces are included regardless of code page type. Possible values are:

UNICODE

Include file spaces that are only in Unicode.

NONUNICODE

Include file spaces that are not only in Unicode.

BOTH

Include file spaces regardless of code page type.

FOLLOWLinks

Specifies whether to display only the files that are stored on the volume or only files that are linked to the volume. You can also display both stored files and linked files. The default is NO. Possible values are:

No

Display only the files that are stored in the volume. Do not display files that have links to the volume.

Yes

Display all files, including files that are stored on the volume and any files that have links to the volume.

JUSTLinks

Display only the files that have links to the volume. Do not display files that are stored on the volume.

Example: Display the contents of a volume for a specific client node

Query the contents of a volume and limit the results to files backed up from the PEGASUS client node.

AIX | **Linux** For the volume /tsmstg/diskvol1.dsm, issue the command:

```
query content /tsmstg/diskvol1.dsm node=pegasus
type=backup
```

Windows For the volume f:\tsmstg\diskvol1.dsm, issue the command:

```
query content f:\tsmstg\diskvol1.dsm node=pegasus
type=backup
```

Results of the command include all logical files that make up any aggregate that is on the volume, even if the aggregate is stored on more than this volume. For aggregates, the query does not determine which logical files are actually stored on the volume for which the query is performed.

Node Name	Type	Filespace Name	FSID	Client's Name for File
PEGASUS	Bkup	\\pegasus\e\$	1	\UNI_TEST\ SM01.DAT
PEGASUS	Bkup	\\pegasus\e\$	1	\UNI_TEST\ SM02.DAT

See Field descriptions for field descriptions.

Example: Display detailed information for a tape volume

Query the contents of the tape volume named WPD001. Display only files that are backed up by the node MARK, and files that are either stored on the volume or linked to the volume. Display only the first four files on the volume.

```
query content wpd001 node=mark count=4 type=backup followlinks=yes
format=detailed
```

```

Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM01.DAT
Hexadecimal Client's Name for File:
Aggregated?: 1/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked: No
Fragment Number:
```

```

Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM02.DAT
Hexadecimal Client's Name for File:
Aggregated?: 2/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked: No
Fragment Number: 2
```

```

Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM03.DAT
Hexadecimal Client's Name for File:
Aggregated?: 3/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked: No
Fragment Number: 3
```

See Field descriptions for field descriptions.

Field descriptions

Node Name

The node to which the file belongs.

Type

The type of file: archive (Arch), backup (Bkup), or space-managed (SpMg) by an IBM Spectrum Protect for Space Management client.

Filespace Name

The file space to which the file belongs.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Hexadecimal Filespace Name

The file space to which the file belongs. If the file space name is in Unicode, the name is displayed in hexadecimal format.

FSID

The file space ID (FSID) for the file space. The server assigns a unique FSID when a file space is first stored on the server.

Client's Name for File

The client's name for the file.

File space names and file names that can be in a different code page or locale than the server do not display correctly in the Operations Center or the administrative command-line interface. The data itself is backed up and can be restored properly, but the file space or file name might display with a combination of invalid characters or blank spaces. The results of the conversion for characters that are not supported by the current code page depends on the operating system. For names that IBM Spectrum Protect is able to partially convert, you might see question marks (??), blanks, unprintable characters, or "...". These characters indicate to the administrator that files do exist.

Hexadecimal Client's Name for File

The client's name for the file that is displayed in hexadecimal format.

Aggregated?

Whether the file is a logical file that is stored as part of an aggregate. If the file is part of an aggregate, the sequence of this file within the aggregate and the total number of logical files in the aggregate are displayed. Results of the command include all logical files that make up any aggregate that is on the volume, even if the aggregate is stored on more than this volume. The query does not determine which logical files are actually stored on the volume for which the query is performed.

If the file is not part of an aggregate, the field displays "no".

Stored Size

The size of the physical file, in bytes. If the file is a logical file that is stored as part of an aggregate, this value indicates the size of the entire aggregate.

Segment Number

For volumes in sequential-access storage pools, specifies whether the physical file (either a single logical file or an aggregate of logical files) is stored across multiple volumes. For example, if the logical file is stored in an aggregate that spans two volumes, the segment number indicates 1/2 (the first part of the physical file is stored on the volume) or 2/2 (the second part of the physical file is stored on the volume). If the segment number is 1/1, the physical file is completely stored on the volume. For volumes in random-access storage pools, no value is displayed for this field.

Cached Copy?

Whether the physical file is a cached copy of a file migrated to the next storage pool. If the file is part of an aggregate, this value pertains to the aggregate.

Linked

Indicates whether the file is stored on the volume or whether the file is linked to the volume.

Fragment Number

Specifies the fragment number. If the fragment number is blank, it is either the first fragment or not a fragment.

Related commands

Table 1. Commands related to QUERY CONTENT

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.

Command	Description
COPY ACTIVE DATA	Copies active backup data.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

AIX Linux Windows

QUERY CONVERSION (Query conversion status of a storage pool)

Use this command to display information about a conversion operation. You can convert a primary storage pool that uses a FILE type device class or a virtual tape library (VTL) to a directory-container storage pool.

Privilege class

To issue this command, you must have restricted storage privilege.

Syntax

```
>>-Query CONVERSION--+-+-----+----->
                        '-pool_name-'

.-Format-----Standard-----
>--+-----+----->>
  '-Format-----+Standard+-'
                        '-Detailed-'
```

Parameters

pool_name

Specifies the source storage pool to query. This parameter is optional. If you do not specify a value for this parameter, information is displayed for all storage pools.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Specify one of the following values:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display conversion information for all storage pools

Display conversion information for all storage pools. See Field descriptions for field descriptions.

```
query conversion
```

Source Storage Pool	Target Storage Pool	Starting Amount	Total Converted	Last Converted
FILEPOOL	CTR	3 GB	3 GB	3 GB
FPOOL	CTR	333 MB	333 MB	267 MB

Example: Display detailed about storage pool conversion

Display detailed information about storage pool conversion. See Field descriptions for field descriptions.

```
query conversion format=detailed

Source Storage Pool: FILEPOOL
Target Storage Pool: CTR
Maximum Processes: 4
    Duration: 60 minutes
Starting Amount: 333 MB
Total Converted: 333 MB
Last Converted: 333 MB
Start Date/Time: 03/24/2016 13:22:32
```

Field descriptions

Source Storage Pool

The name of the storage pool that is being converted.

Target Storage Pool

The name of the destination storage pool, where the converted data will be stored.

Maximum Processes

Specifies the maximum number of conversion processes.

Duration

Specifies the length of time, in minutes, for conversion.

Starting Amount

The starting amount of data to convert, in megabytes (MB), gigabytes (GB), or terabytes (TB).

Total Converted

The total amount of data that is converted, in megabytes (MB), gigabytes (GB), or terabytes (TB).

Last Converted

The amount of data, in megabytes (MB), gigabytes (GB), or terabytes (TB), that is converted during this conversion process.

Start Date/Time

The time and date that the CONVERT STGPOOL command was first issued for the storage pool.

Related commands

Table 1. Commands related to QUERY CONVERSION

Command	Description
CONVERT STGPOOL	Convert a storage pool to a directory-container storage pool.
QUERY CLEANUP	Query the cleanup status of a source storage pool.

QUERY COPYGROUP (Query copy groups)

Use this command to display information about one or more copy groups.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query Copygroup----->
. -*--*--*--STANDARD-----
>-----+-----+----->
|          .-*--*--STANDARD-----|
|'-domain_name-----+-----|
|          |          .-*--STANDARD-----| |
|          |'-policy_set_name-----+-----|
|          |          |          .-STANDARD-. |
|          |          |          '-class_name-----+-----|
|          |          |          '-STANDARD-'
|          |          |
.-Type-----Backup----- .-Format-----Standard-----.
```

```
>-----<
'-Type-----+Backup--+' '-Format-----+Standard-+-'
          '-Archive-'          '-Detailed-'
```

Parameters

domain_name

Specifies the policy domain that is associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named copy group.

policy_set_name

Specifies the policy set associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy sets are queried. You must specify this parameter when querying an explicitly named copy group.

class_name

Specifies the management class that is associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all management classes are queried. You must specify this parameter when querying an explicitly named copy group.

STANDARD

Specifies the name of the copy group. This parameter is optional. The name of the copy group must be STANDARD. The default is STANDARD.

Type

Specifies the type of copy group to be queried. This parameter is optional. The default value is BACKUP. Possible values are:

Backup

Specifies that you want to query backup copy groups.

Archive

Specifies that you want to query archive copy groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display information about the default backup copy group

Display information about the default backup copy group in the ENGPOLDOM engineering policy domain. Issue the following command:

```
query copygroup engpoldom * *
```

The following data shows the output from the query. It shows that the ACTIVE policy set contains two backup copy groups that belong to the MCENG and STANDARD management classes.

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
ENGPOLDOM	ACTIVE	MCENG	STANDARD	5	4	90	600
ENGPOLDOM	ACTIVE	STANDARD	STANDARD	2	1	30	60
ENGPOLDOM	STANDARD	MCENG	STANDARD	5	4	90	600
ENGPOLDOM	STANDARD	STANDARD	STANDARD	2	1	30	60
ENGPOLDOM	TEST	STANDARD	STANDARD	2	1	30	60

Example: Display detailed information on one backup copy group

Display complete information on the backup copy group assigned to the ACTIVEFILES management class in the VACATION policy set of the EMPLOYEE_RECORDS policy domain. Issue the command:

```
query copygroup employee_records vacation
activefiles format=detailed
```

Example: Display information on the backup copy group in the STANDARD management class and policy set

From a managed server, display complete information on the backup copy group assigned to the STANDARD management class in the STANDARD policy set of the ADMIN_RECORDS policy domain. Issue the command:

```
query copygroup admin_records
standard standard format=detailed

Policy Domain Name: ADMIN_RECORDS
Policy Set Name: STANDARD
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 2
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: BACKUPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2002.10.02 17.51.49
Managing profile: ADMIN_INFO
Changes Pending: Yes
```

Example: Display information on an archive copy group

From a managed server, display complete information on the archive copy group STANDARD that is assigned to the MCLASS1 management class in the SUMMER policy set of the PROG1 policy domain. Issue the command:

```
query copygroup prog1 summer mclass1
type=archive format=detailed

Policy Domain Name: PROG1
Policy Set Name: SUMMER
Mgmt Class Name: MCLASS1
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 730
Retention Initiation: Creation
Minimum Retention:
Copy Serialization: Shared Static
Copy Frequency: Cmd
Copy Mode: Absolute
Copy Destination: ARCHPOOL
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2002.10.02 17.42.49
Managing profile: ADMIN_INFO
```

Example: Display information on the copy group for a NAS backup

Query the copy group for the NAS backup. Issue the command:

```
query copygroup nasdomain
type=backup

Policy Domain Name: NASDOMAIN
Policy Set Name: ACTIVE
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 2
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
```


Copy Destination: NASPOOL
 Table of Contents (TOC) Destination: BACKUPPOOL
 Last Update by (administrator): SERVER_CONSOLE
 Last Update Date/Time: 10/02/2002 12:16:52
 Managing profile:
 Changes Pending: Yes

Field descriptions

Policy Domain Name
 The name of the policy domain.

Policy Set Name
 The name of the policy set.

Mgmt Class Name
 The name of the management class.

Copy Group Name
 The name of the copy group. This name is always STANDARD.

Copy Group Type
 The type of the copy group.

Versions Data Exists
 The maximum number of backup versions to retain for files that are currently on the client file system.

Versions Data Deleted
 The maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using IBM Spectrum Protect™.

Retain Extra Versions
 The number of days to retain a backup version after that version becomes inactive.

Retain Only Version
 The number of days to retain the last backup version of a file that has been deleted from the client file system.

Copy Serialization
 Whether a file can be in use during an archive operation.

Copy Frequency
 The copy frequency of the copy group. For archive copy groups, this value is always CMD.

Copy Mode
 Specifies that files in the copy group are archived regardless of whether they have been modified. For archive copy groups, this value is always ABSOLUTE.

Copy Destination
 The name of the storage pool where the server initially stores files associated with this archive copy group.

Table of Contents (TOC) Destination
 The name of the primary storage pool in which TOCs are initially stored for image backup operations in which TOC generation is requested.

Last Update by (administrator)
 The name of the administrator or server that most recently updated the copy group. If this field contains \$\$CONFIG_MANAGER\$\$, the copy group is associated with a domain that is managed by the configuration manager.

Last Update Date/Time
 The date and time when the copy group was most recently defined or updated.

Managing Profile
 The profile or profiles to which the managed server subscribed to get the definition of this policy copy group.

Changes Pending
 Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 1. Commands related to QUERY COPYGROUP

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

AIX | Linux | Windows

QUERY DAMAGED (Query damaged data in a directory-container or cloud-container storage pool)

Use this command to display information about damaged data extents in a directory-container or cloud-container storage pool. Use this command together with the AUDIT CONTAINER command to determine a recovery method for the damaged data.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query DAMaged--pool_name----->
.-Type----Status-----
>--+-----+-----><
'-Type-----+INventory-----+'
      +-Node--| A |-----+
      '-CONTAINER--| A |-'
```

A (Additional filter by node name)

```
|--+-----+-----|
'-Nodename-----node_name-'
```

Parameters

pool_name (Required)

Specifies the name of the directory-container or cloud storage pool.

Type

Specifies the type of information to display. This parameter is optional. Specify one of the following values:

Status

Specifies that information is displayed about damaged data extents. For cloud storage pools, orphaned extents are also displayed. This is the default.

Node

Specifies that information about the number of damaged files per node is displayed.

INventory

Specifies that inventory information for each damaged file is displayed.

CONTAINER

Specifies that the containers that contain damaged data extents or cloud orphaned extents are displayed. For directory-container storage pools, storage pool directories are also displayed.

Nodename

Specifies that damaged file information for a single node is displayed.

Restriction: You cannot specify this parameter if the TYPE=CONTAINER or TYPE=STATUS parameter is specified.

Example: Display status information about damaged or orphaned data extents

Display information about the status of damaged data extents that are stored in a container.

```
query damaged pool1 type=status
```

Storage Pool Name	Non-Dedup Data Extent Count	Dedup Data Extent Count	Cloud Orphaned Extent Count
POOL1	58	145	

For cloud storage pools, the number of orphaned extents is also displayed.

Storage Pool Name	Non-Dedup Data Extent Count	Dedup Data Extent Count	Cloud Orphaned Extent Count
-------------------	-----------------------------	-------------------------	-----------------------------

```
-----
POOL1                65                238                18
-----
```

Example: Display information about a damaged file for a node type

Display information about damaged files that are stored in a node.

```
query damaged pool1 type=node
```

```
Node Name      Number of
              Damaged Files
-----
POOL1          37
```

Example: Display information about a damaged file for an inventory type

Display information about damaged files that are stored in an inventory.

```
query damaged pool2 type=inventory
```

```
Client's Name for File: /data/files/10.out
                        Type: Bkup
                        Node Name: NODE1
Filespace Name: /data/space
                        State: Available
Insertion time: 01/19/2015 16:01:35
Object ID: 2073
```

Example: Display information about a damaged file for a container type

Display information about damaged files that are stored in a container.

```
query damaged pool3 type=container
```

```
Directory ID: 1
Directory: /abc/space/container1
Container: /abc/space/container1/00/0000000000000022.dcf
State: Unavailable
```

For cloud containers, only the name of the container is displayed.

```
Directory ID:
Directory:
Container: ibmsp.12520ae05b4011e613320a0027000000/
          001-10006a3278bc34f0e4118a850090fa3dcb48/
          000000000000001.ncf
State:
```

For local storage, the following information about a damaged container is displayed.

```
Directory ID: 1
Directory: localdirectory
Container: localdirectory/00/0000000000000011.ncf
State: Unavailable
```

Field descriptions

Client's Name for File (TYPE=INVENTORY only)

The name of the file.

Cloud Orphaned Extent Count (TYPE=STATUS only)

The number of orphaned extents in a cloud storage pool. Extents are considered orphaned if they do not have a corresponding database entry.

Container (TYPE=CONTAINER only)

The name of the container.

Deduplicated Extent Count (TYPE=STATUS only)

The number of damaged extents in the storage pool for deduplicated data.

Directory (TYPE=CONTAINER only)

The name of the storage pool directory.
 Directory ID (TYPE=CONTAINER only)
 The identification number of the storage pool directory.

Filespace Name (TYPE=INVENTORY only)
 The name of file space.

Insertion time (TYPE=INVENTORY only)
 The date and time that the object was stored on the server.

Node Name (TYPE=INVENTORY or TYPE=NODE only)
 The name of the node.

Non-Deduplicated Extent Count (TYPE=STATUS only)
 The number of damaged extents in the storage pool for data that is not deduplicated, such as metadata and client-encrypted data.

Number of Damaged Files (TYPE=NODE only)
 The number of damaged files per node.

Object ID (TYPE=INVENTORY only)
 The identification number of the object.

State (TYPE=INVENTORY or TYPE=CONTAINER only)
 The state of the data in either the inventory or the container, depending on the type of data you are querying. The field can contain one of the following values:

- Active
The version of the file in the inventory is active. There can be only one active version of the file in the inventory.
- Inactive
The version of the file in the inventory is inactive. There can be multiple inactive versions of the file in the inventory.
- Available
The state of the container is available.
- Unavailable
The state of the container is unavailable. For example, a container might be unavailable if the header is corrupted or if the container cannot be opened.
- Read-Only
The container is in a read-only state. Data in the container can be read, but data cannot be written to the container.
- Pending
The container is pending deletion. The contents of the container were moved to a different container, and the container is ready to be deleted.

Type (TYPE=INVENTORY only)
 The type of data in the file.

Table 1. Commands related to QUERY DAMAGED

Command	Description
AUDIT CONTAINER	Audit a directory-container storage pool.
QUERY CLEANUP	Query the cleanup status of a source storage pool.
QUERY CONTAINER	Displays information about a container.
REMOVE DAMAGED	Removes damaged data from a source storage pool.

QUERY DATAMOVER (Display data mover definitions)

Use this command to display data mover definitions.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query DATAMover-----*-----
      '-data_mover_name-'

.-Format-----Standard-----
>-----+-----+-----+-----+----->
      '-Format-----+Standard-+-'
              '-Detailed-'

.-Type-----*-----
>-----+-----+-----+-----+----->>
      |                                     (1) (2) |
      '-Type-----+NAS-----+-----+'
              +-NASCLUSTER-+
              '-NASVSERVER-'

```

Notes:

1. You must specify the TYPE parameter if FORMAT=DETAILED.
2. You can specify TYPE=NASCLUSTER and TYPE=NASVSERVER only on an AIX, Linux, or Windows operating system.

Parameters

data_mover_name

Specifies the name of the data mover to display. You can specify multiple names with a wildcard character. The default displays all data movers.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD.

Standard

Specifies that name and address information is displayed.

Detailed

Specifies that complete information is displayed.

Type

Specifies the type of data mover to be displayed. If you specify FORMAT=DETAILED, you must specify a value for the TYPE parameter.

NAS

Specifies a NAS file server.

AIX	Linux	Windows	NASCLUSTER
-----	-------	---------	------------

AIX	Linux	Windows	Specifies a clustered NAS file server.
-----	-------	---------	--

AIX	Linux	Windows	NASVSERVER
-----	-------	---------	------------

AIX	Linux	Windows	Specifies a virtual storage device within a cluster.
-----	-------	---------	--

Example: Display information about all data movers

Display the data movers on the server. Issue the command:

```
query datamover
```

Data Mover Name	Data Mover Type	Online
NASMOVER1	NAS	Yes
NASMOVER2	NAS	No

See Field descriptions for field descriptions.

Example: Display information about one data mover

Display partial information about data mover DATAMOVER6. Issue the command:

```
query datamover datamover6 type=nas
```

Source Name	Type	Online
-----	-----	-----

DATAMOVER6 NAS Yes

See Field descriptions for field descriptions.

Example: Display detailed information about one data mover

Display detailed information about data mover DATAMOVER6. The TYPE parameter is required when FORMAT=DETAILED. Issue the command:

```
query datamover datamover6 format=detailed type=nas

      Data Mover Name:   DataMover6
      Data Mover Type:   NAS
      IP Address:        198.51.100.0
      TCP/IP Port Number: 10000
      User Name:         NDMPadmin
      Storage Pool Data Format: NDMPDUMP
      Online:            Yes
      Last Update by (administrator): ADMIN
      Last Update Date/Time: 05/23/2015 09:26:33
```

See Field descriptions for field descriptions.

[AIX](#) [Linux](#) [Windows](#)

Example: Display detailed information about a clustered NAS data mover

Display detailed information about a clustered NAS data mover that is named CLUSTERA. Issue the following command:

```
query datamover clustera format=detailed type=nascluster

      Data Mover Name:   CLUSTERA
      Data Mover Type:   NASCLUSTER
      IP Address:        192.0.2.255
      TCP/IP Port Number: 10000
      User Name:         ndmp
      Storage Pool Data Format: NETAPPDUMP
      Online:            Yes
      Last Update by (administrator): ADMIN
      Last Update Date/Time: 04/28/2015 09:26:33
```

See Field descriptions for field descriptions.

Field descriptions

Data Mover Name

Specifies the name of the data mover.

Data Mover Type

Specifies the type of the data mover.

IP Address

Specifies the IP address of the data mover.

TCP/IP Port Number

Specifies the TCP port number for the data mover.

User Name

Specifies the user ID that the server uses to access the data mover.

Storage Pool Data Format

Specifies the data format that is used by the data mover.

Online

Specifies whether the data mover is online and available for use.

Last Update by (administrator)

Specifies the ID of the administrator who completed the last update.

Last Update Date/Time

Specifies the date and time when the last update occurred.

Related commands

Table 1. Commands related to QUERY DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DELETE DATAMOVER	Deletes a data mover.
UPDATE DATAMOVER	Changes the definition for a data mover.

QUERY DB (Display database information)

Use this command to display information about the database.

Privilege class

Any administrator can issue this command.

Syntax

```

.-Format-----Standard-----
>>-Query DB--+-----+----->>
      '-Format-----+Standard--'
              '-Detailed-'

```

Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. The following values are possible:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary statistics about the database

Display statistical information about the database. Issue the command:

```
query db
```

Database Name	Total Pages	Usable Pages	Used Pages	Free Pages
TSMDB1	32,776	32,504	24,220	8,284

See Field descriptions for field descriptions.

Example: Display detailed database information

Display detailed statistical information about the database. Issue the command:

```
query db format=detailed
```

```

Database Name: TSM_DB2
Total Space of File System (MB): 1,748,800
Space Used on File System (MB): 2,304,355
Space Used by Database (MB): 448
Free Space Available (MB): 235,609
  Total Pages: 32,776
  Usable Pages: 32,504
  Used Pages: 24,220
  Free Pages: 8,284
Buffer Pool Hit Ratio: 99.3
Total Buffer Requests: 204,121
Sort Overflows: 0
Package Cache Hit Ratio: 89.8

```

Last Database Reorganization: 05/25/2009 16:44:06
Full Device Class Name: FILE
Number of Database Backup Streams: 4
Incrementals Since Last Full: 0
Last Complete Backup Date/Time: 05/18/2009 22:55:19
Compress Database Backups: Yes
Protect Master Encryption Key: No

See Field descriptions for field descriptions.

Field descriptions

Database Name

The name of the database that is defined and configured for use by the IBM Spectrum Protect™ server.

AIX | **Linux** Total Space of File System (MB)

The total space, in megabytes, of the file systems in which the database is located.

Windows Total Space of File System (MB)

The total space, in megabytes, of the drives on which the database is located.

Space Used on File System (MB)

The amount of database space, in megabytes, that is in use.

Space Used by Database (MB)

The size of the database, in megabytes. The value does not include any temporary table space. The size of the database is calculated from the amount of space that is used on the file system containing the database.

Free Space Available (MB)

The amount of database space, in megabytes, that is not in use.

Total Pages

The total number of pages in the table space.

Usable Pages

The number of usable pages in the table space.

Used Pages

The number of used pages in the table space.

Free Pages

The total number of free pages in all table spaces. The IBM Spectrum Protect database has up to 10 table spaces.

Buffer Pool Hit Ratio

The total hit ratio percent.

Total Buffer Requests

The total number of buffer pool data logical reads and index logical reads since the last time the database was started or since the database monitor was reset.

Sort Overflows

The total number of sorts that ran out of the sort heap and might have required disk space for temporary storage.

Package Cache Hit Ratio

A percentage that indicates how well the package cache is helping to avoid reloading packages and sections for static SQL from the system catalogs. It also indicates how well the package cache is helping to avoid recompiling dynamic SQL statements. A high ratio indicates that it is successful in avoiding these activities.

Last Database Reorganization

The last time that the database manager completed an automatic reorganization activity.

Full Device Class Name

The name of the device class that is used for full database backups.

Number of Database Backup Streams

The number of concurrent data movement streams that were used during the database backup.

Incrementals Since Last Full

The number of incremental backups that were completed since the last full backup.

Last Complete Backup Date/Time

The date and time of the last full backup.

Compress Database Backups

Specifies whether database backups are compressed.

Protect Master Encryption Key

Specifies whether database backups include a copy of the server master encryption key.

Related commands

Table 1. Commands related to QUERY DB

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
EXTEND DBSPACE	Adds directories to increase space for use by the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

QUERY DBSPACE (Display database storage space)

Use this command to display information about the directories used by the database to store data.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-QUERY DBSpace-----<<
```

Parameters

None.

Example: Display database storage space information

Display information about database storage space. Issue the command:

```
query dbspace
```

AIX		Linux	
Location	Total Space of File System (MB)	Used Space on File System (MB)	Free Space Available (MB)
/tsmdb001	1,748,800	1,513,191.125	117,804.422
/tsmdb002	1,748,800	1,513,191.125	117,804.422

Windows			
Location	Total Space of File System (MB)	Used Space on File System (MB)	Free Space Available (MB)
d:\tsm\db001	1,748,800	1,513,191.125	117,804.422
e:\tsm\db002	1,748,800	1,513,191.125	117,804.422

See Field descriptions for field descriptions.

Field descriptions

Location

Specifies the locations of database directories.

AIX Total Space of File System (MB)

AIX The total amount of space, in megabytes, of the file system in which the database is located.

Windows Total Space of File System (MB)

Windows The total amount of space, in megabytes, of the drives on which the database is located.

Used Space on File System (MB)

The amount of storage space, in megabytes, that is in use.

AIX **Linux** When you run the QUERY DBSPACE command, the value in the output might be greater than the value that is obtained by running the df system command. The output from the df system command does not include the amount of space that is reserved for the root user.

Linux If you run the `df` system command, the default percentage of space that is reserved for the root user is 5%. You can change this default value.

Free Space Available (MB)

The amount of space, in megabytes, that is not in use.

Windows Free Space Available (MB)

Windows The amount of space remaining on the drive where the directory is located.

Related commands

Table 1. Commands related to QUERY DBSPACE

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
EXTEND DBSPACE	Adds directories to increase space for use by the database.
QUERY DB	Displays allocation information about the database.

AIX Linux Windows

QUERY DEDUPSTATS (Query data deduplication statistics)

Use this command to display information about data deduplication statistics for a directory-container storage pool or a cloud storage pool. You can display statistics for an entire storage pool or for data from a specified group of client nodes.

You must issue the `GENERATE DEDUPSTATS` command before you can issue the `QUERY DEDUPSTATS` command.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query DEDUPStats--+-+-----+----->
                        '-pool_name-'

.-,------.
v          | .-*------.
>-----+-----+-----+----->
  '+-node_name-----+' | .-,------. |
  '-node_group_name-'  | v          | |
                        +---+-----+---+
                        | .-,------. |
                        | v          | |
                        +-----FSID-----'

.-Format-----Standard-----.
>-----+-----+-----+----->
  '-Format-----+Standard+-'
                +-Detailed++
                '-SUMmary--'

.-CODEType-----BOTH------.
>-----+-----+-----+----->
  '-CODEType-----+UNICODE-----+'
                +-NONUNICODE++
                '-BOTH-----'

.-NAMEType-----SERVER------.
>-----+-----+-----+-----+----->
  '-NAMEType-----+SERVER--+-' '-BEGINDate-----date-'
                +-UNICODE++
                '-FSID-----'

>-----+-----+-----+-----+----->
  '-BEGINTime-----time-' '-ENDDate-----date-'
```


Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for file spaces that are in Unicode format. You can use this parameter for IBM Spectrum Protect™ clients that use Windows, NetWare, or Macintosh OS X operating systems.

Use this parameter only when you enter a node name and a file space name or FSID.

Restriction: When you specify this parameter, the file space name cannot contain a wildcard.

Specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names. This is the default.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their FSIDs.

BEGINDate

Specifies the start date to query data deduplication statistics. This parameter is optional. You can use this parameter with the BEGINTIME parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time is at 12 midnight on the date you specify.

Restriction: You can specify this parameter only when you specify the ALLSTATS=YES parameter.

Specify one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date.	09/15/2015
TODAY	The current date.	TODAY
TODAY-days or days	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include records that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include records that were active on the 10th day of the current month.

BEGINTime

Specifies the start time to query the data deduplication statistics. This parameter is optional. You can use this parameter with the BEGINDATE parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date is the current date at the time you specify.

Restriction: You can specify this parameter only when you specify the ALLSTATS=YES parameter.

Specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	10:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00.

ENDDate

Specifies the end date to query data deduplication statistics. This parameter is optional. You can use this parameter with the ENDTIME parameter to specify a range for the date and time. If you specify an end date without an end time, the time is

at 11:59:59 p.m. on the specified end date.

Restriction: You can specify this parameter only when you specify the ALLSTATS=YES parameter.

Specify one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include records that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include records that were active on the 10th day of the current month.

ENDTime

Specifies the end time of the range to query the data deduplication statistics. This parameter is optional. You can use this parameter with the ENDDATE parameter to specify a range for the date and time. If you specify an end time without an end date, the date is the current date at the time you specify.

Restriction: You can specify this parameter only when you specify the ALLSTATS=YES parameter.

Specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	10:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+02:00 or +02:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-02:00 or -02:00.

ALLStats

Specifies whether to display all data deduplication statistics or only the most recently generated data deduplication statistics. This parameter is optional. Specify one of the following values:

No

Displays only data deduplication statistics that were most recently generated for each node and file space.

Yes

Displays all data deduplication statistics.

REPortid

Specifies an ID for a set of data deduplication statistics that is generated on a specific day for specified nodes, file spaces, or both. For example, if you generate statistics on 30 September 2018 for a node list (TEST1, TEST2, TEST3, and MYGROUP1) and a file space list (FS1, FS2, and /tmp*), a report ID (for example, 1) is assigned to that set. If statistics are generated for the same nodes and file spaces on the next day, a new report ID (for example, 2) is assigned to that set. This parameter is optional.

DESCription

Specifies a description of the generated statistics. This parameter is optional.

Example: View data deduplication statistics in standard format

Display data deduplication statistics for a storage pool that is named POOL1. The data deduplication statistics are for node NODE1 and the statistics from 8 May 2015 are displayed. See Field descriptions for field descriptions.

```
query dedupstats pool1 node1 begindate=05/08/2015
```

```
Date/Time: 05/05/2015 15:15:23
Storage Pool Name: POOL1
Node Name: NODE1
Filespace Name: \\fs1\al
FSID: 41
Type: Bkup
Total Saving Percentage: 86.62
Total Data Protected (MB): 311
```

Example: View detailed data deduplication statistics

Display detailed information for data deduplication for a storage pool that is named POOL1.

```
query dedupstats pool1 format=detailed
```

```
Date/Time: 05/05/2015 15:15:23
Storage Pool Name: POOL1
Node Name: NODE1
Filespace Name: \\fs1\al
FSID: 41
Type: Bkup
Total Data Protected (MB): 47,646
Total Space Used (MB): 10,139
Total Space Saved (MB): 37,507
Total Saving Percentage: 78.72
Deduplication Savings: 16,228,107,499
Deduplication Percentage: 42.59
Non-Deduplicated Extent Count: 1,658
Non-Deduplicated Extent Space Used: 732,626
Unique Extent Count: 189,791
Unique Extent Space Used: 23,385,014,635
Shared Extent Count: 178,712
Shared Extent Data Protected: 26,575,010,669
Shared Extent Space Used: 5,267,815,421
Compression Savings: 5,267,815,421
Compression Percentage: 62.93
Compressed Extent Count: 352,498
Uncompressed Extent Count: 17,663
Encryption Extent Space Used: 52,901,672
Encryption Percentage: 100.00
Encrypted Extent Count: 188
Unencrypted Extent Count: 0
Report ID: 1
Description:
```

Example: View summarized data deduplication statistics

Display a summary of information for a set of statistics.

```
query dedupstatus reportid=1234 format=summary
```

```
Report ID: 1234
Description:
Date/Time: 09/15/2017 16:59:55
Storage Pool Name: DIRPOOL
Node Name: TEST1,TEST2,TEST3,MYGROUP1
Filespace Name: FS1,FS2,/tmp*
Type: Bkup
Total Data Protected (MB): 47,646
Total Space Used (MB): 10,139
Total Space Saved (MB): 37,507
Total Saving Percentage: 78.72
Deduplication Savings: 16,228,107,499
Deduplication Percentage: 42.59
Non-Deduplicated Extent Count: 1,658
Non-Deduplicated Extent Space Used: 732,626
Unique Extent Count: 189,791
Unique Extent Space Used: 23,385,014,635
Shared Extent Count: 178,712
Shared Extent Data Protected: 26,575,010,669
Shared Extent Space Used: 5,267,815,421
Compression Savings: 5,267,815,421
```

Compression Percentage: 62.93
Compressed Extent Count: 352,498
Uncompressed Extent Count: 17,663
Encryption Extent Space Used: 52,901,672
Encryption Percentage: 100.00
Encrypted Extent Count: 188
Unencrypted Extent Count: 0

Field descriptions

Report ID

An ID for a set of data deduplication statistics that is generated on a specific day for a specified group of nodes, file spaces, or both.

Description

A description of the statistics set that is generated.

Date/Time

The time and date that the data deduplication statistics are generated.

Storage Pool Name

The name of the storage pool.

Node Name

The name of the client node whose data is contained in the data deduplication statistics.

Filespace Name

The name of the file space.

FSID

The name of the file space identifier.

Type

The type of data. The following values are possible:

Arch

Data that is archived.

Bkup

Data that is backed up.

SpMg

Data that is migrated from an IBM Spectrum Protect for Space Management client.

Total Data Protected (MB)

The logical amount of data, in megabytes, that is protected in the storage pool before data deduplication and compression. This value represents the sum of the Total Space Used (MB) and Total Space Saved (MB) values.

Total Space Used (MB)

The total amount of used space in the storage pool, in megabytes. This value is the physical amount of data that is backed up after data deduplication and compression.

Total Space Saved (MB)

The total amount of space, in megabytes, of data that is removed from the storage pool because of data deduplication and compression. This value represents the sum of the Deduplication Savings and Compression Savings values.

Total Saving Percentage

The percentage of data that is removed from the storage pool because of compression and data deduplication.

Deduplication Savings

The amount of used space that is saved in the storage pool because of data deduplication.

Deduplication Percentage

The percentage of data that is removed from the storage pool because of data deduplication.

Non-Deduplicated Extent Count

The number of data extents that are not deduplicated in the storage pool.

Non-Deduplicated Extent Space Used

The amount of space that is used by data extents that are not deduplicated in the storage pool. This value applies to containers that have a .ncf file type and that do not have deduplicated data.

Tip: Data extents that are not deduplicated consist of the following data or file types:

- File metadata.
- Files that are less than 2 KB.
- Files that use client encryption.

Unique Extent Count

The number of data extents that are not shared by a node.

Unique Extent Space Used
The amount of space in the storage pool that is not shared by a node. This value applies to containers that have a .dcf file type and that do not have deduplicated data.

Shared Extent Count
The number of data extents that are used multiple times by the same node or by different nodes because of data deduplication.

Shared Extent Data Protected
The amount of space in the storage pool that is protected by shared data extents before data deduplication.

Shared Extent Space Used
The amount of space in the storage pool that is used by shared data extents after data deduplication.

Compression Savings
The amount of used space that is saved in the storage pool because of compression after data deduplication.

Compression Percentage
The percentage of data that is removed from the storage pool because of compression.

Compressed Extent Count
The number of data extents that are compressed.

Uncompressed Extent Count
The number of data extents that are uncompressed.

Encryption Extent Space Used
The amount of space in the storage pool that is used by encrypted data extents.

Encryption Percentage
The percentage of encrypted data in the storage pool.

Encrypted Extent Count
The number of data extents that are encrypted.

Unencrypted Extent Count
The number of data extents that are not encrypted.

Related commands

Table 1. Commands related to QUERY DEDUPSTATS

Command	Description
DELETE DEDUPSTATS	Deletes data deduplication statistics.
GENERATE DEDUPSTATS	Generates data deduplication statistics.

QUERY DEVCLASS (Display information on one or more device classes)

Use this command to display information on one or more device classes.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query DEVclass-.*-----+----->
                    +-----+----->
                    '-device_class_name-'

.-Format----Standard----.
>--+-----+----->>
  '-Format----Standard--+'
                    '-Detailed-'

```

Parameters

device_class_name
Specifies the name of the device class to be queried. This parameter is optional. You can use wildcard characters to specify this name. All matching device classes are displayed. If you do not specify a value for this parameter, all device classes are

displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified device class.

Detailed

Specifies that complete information is displayed for the specified device class.

Example: List all device classes

Display information on all device classes.

```
query devclass
```

AIX	Linux	Windows	Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
			8MMTAPE	Sequential	1	8MM	DRIVE	6,144.0	2
			DISK	Random	4				
			PLAINFILES	Sequential	1	FILE		50.0	1
			8MMSP2	Sequential	2	8MM	DRIVE	44.4	DRIVES

See Field descriptions for field descriptions.

Example: Display detailed information for a specific FILE device class

Display information in full detail on the PLAINFILES device class.

```
query devclass plainfiles format=detailed
```

```

Device Class Name: PLAINFILES
Device Access Strategy: Sequential
Storage Pool Count: 1
Device Type: FILE
Format:
Est/Max Capacity (MB): 50.0
Mount Limit: 1
Mount Wait (min):
Mount Retention (min):
Label Prefix:
Drive Letter:
Library:
Directory:
Server Name:
Retry Period:
Retry Interval:
Shared:
Primary Allocation (MB):
Secondary Allocation (MB):
Compression:
Retention:
Protection:
Expiration Date:
Unit:
Logical Block Protection:
Last Update by (administrator): ADMIN
Last Update Date/Time: 05/31/2000 13:15:36

```

See Field descriptions for field descriptions.

Example: Display detailed information for a specific 3592 device class

Display full details on the 3592 device class.

```
query devclass 3592 format=detailed
```

```

Device Class Name: 3592
Device Access Strategy: Sequential
Storage Pool Count: 1
Device Type: 3592
Format: 3592
Est/Max Capacity (MB):
Mount Limit: DRIVES
Mount Wait (min): 60
Mount Retention (min): 60
Label Prefix: ADSM
Windows Drive Letter:
Library: MANLIB
Directory:
Server Name:
Retry Period:
Retry Interval:
AIX Linux Windows Shared:
High-level Address:
WORM: No
Scaled Capacity: 90
Drive Encryption: On
AIX Linux Primary Allocation (MB):
Secondary Allocation (MB):
Compression:
Retention:
Protection:
Expiration Date:
Unit:
Logical Block Protection: Read/Write
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 08/04/03 14:28:31

```

See Field descriptions for field descriptions.

Field descriptions

Device Class Name

The name of the device class.

Device Access Strategy

How data is written to the device class.

Storage Pool Count

The number of storage pools that are assigned to the device class.

Device Type

The device type of the device class.

Format

The recording format.

Est/Max Capacity (MB)

The estimated or maximum capacity of a volume that is associated with the device class.

Mount Limit

The maximum number of sequential access volumes that can be mounted concurrently or specifies that DRIVES is the mount limit.

Mount Wait (min)

The maximum number of minutes to wait for a sequential access volume to be mounted.

Mount Retention (min)

The number of minutes to retain an idle sequential access volume before dismounting it.

Label Prefix

The high-level qualifier of the data set name that the server writes into the sequential access media labels.

Windows Drive Letter

The drive letter for a removable file.

Library

The name of the defined library object that contains the drives that are used by the device class.

Directory

The directory or directories for a shared FILE device class.

Server Name

The name of a defined server.

Retry Period

The interval over which the server attempts to contact a target server if communications failure is suspected.

Retry Interval

How often the retries are done within a retry period.

Shared

Whether this FILE device class is shared between the server and one or more storage agents.

High-level Address

The IP address of the device in dotted decimal format.

Minimum Capacity

The minimum capacity of a volume that is associated with the device class.

WORM

Whether this drive is a write once, read many (WORM) device.

Drive Encryption

Whether drive encryption is allowed. This field applies only to volumes in a storage pool that is associated with a device type of 3592, LTO, or ECARTRIDGE.

Scaled Capacity

The percentage of the media capacity that can be used to store data.

AIX | **Linux** Primary Allocation (MB)

AIX | **Linux** For FILE device classes that represent storage that is managed by a z/OS® media server. Specifies the initial amount of space that is dynamically allocated when a new volume is opened.

AIX | **Linux** Secondary Allocation (MB)

AIX | **Linux** For FILE device classes that represent storage that is managed by a z/OS media server. Specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up.

AIX | **Linux** Compression

AIX | **Linux** For tape device classes that represent storage that is managed by a z/OS media server. Specifies whether the data is compressed.

AIX | **Linux** Retention

AIX | **Linux** For tape device classes that represent storage that is managed by a z/OS media server. Specifies the number of days to retain the tape, if retention is used.

AIX | **Linux** Protection

AIX | **Linux** For tape device classes that represent storage that is managed by a z/OS media server. Specifies whether the volumes are protected by the RACF program.

AIX | **Linux** Expiration Date

AIX | **Linux** For tape device classes that represent storage that is managed by a z/OS media server. Specifies the expiration date that is placed on the tape labels for this device class, if expiration is used.

AIX | **Linux** Unit

AIX | **Linux** For tape device classes that represent storage that is managed by a z/OS media server. Specifies the esoteric unit name for the group of tape devices.

Logical Block Protection

Specifies whether logical block protection is enabled and, if it is, the mode. Possible values are Read/Write, Write-only, and No. You can use logical block protection only with the following types of drives and media:

- IBM® LTO5 and later
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later
- Oracle StorageTek T10000C and T10000D drives

Last Update by (administrator)

The administrator that made the last update to the device class.

Last Update Date/Time

The date and time of the last update.

Related commands

Table 1. Commands related to QUERY DEVCLASS

Command	Description
DEFINE DEVCLASS	Defines a device class.
AIX Linux DEFINE DEVCLASS (z/OS media server)	AIX Linux Defines a device class to use storage managed by a z/OS media server.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DEVCLASS	Deletes a device class.
QUERY DIRSPACE	Displays information about FILE directories.

Command	Description
QUERY SERVER	Displays information about servers.
UPDATE DEVCLASS	Changes the attributes of a device class.
AIX Linux UPDATE DEVCLASS (z/OS media server)	AIX Linux Changes the attributes of a device class for storage managed by a z/OS media server.

QUERY DIRSPACE (Query storage utilization of FILE directories)

Use this command to display information about free space in the directories associated with a device class with a device type of FILE.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query DIRSPace--+-+-----+-----<<
                    '-device_class_name-'
```

Parameters

device_class_name

Specifies the name of the device class to be queried. This parameter is optional. You can use wildcard characters to specify this name. All matching device classes of device type FILE are displayed. If you do not specify a value for this parameter, all device classes of device type FILE are displayed.

Example: List FILE type device classes

Display information for all device classes with a device type of FILE. In the following example the unit M is equivalent to megabytes, and the unit G is equivalent to gigabytes.

```
query dirspace
```

Windows

Device Class	Directory	Estimated Capacity	Estimated Available
DBBKUP	/This/is/a/large/directory	13,000 M	5,543 M
DBBKUP	/This/is/directory2	13,000 M	7,123 M
DBBKUP2	/This/is/a/huge/directory	2,256 G	2,200 G

Windows

Device Class	Directory	Estimated Capacity	Estimated Available
DBBKUP	G:\This\is\a\large\directory	13,000 M	5,543 M
DBBKUP	G:\This\is\directory2	13,000 M	7,123 M
DBBKUP2	G:\This\is\a\huge\directory	2,256 G	2,200 G

Field descriptions

Device Class Name

The name of the device class.

Directory

The path of the directory located on the server.

Estimated Capacity

The estimated total capacity for the directory.

Estimated Available

The estimated remaining available space for the directory.

Related commands

Table 1. Commands related to QUERY DIRSPACE

Command	Description
DEFINE DEVCLASS	Defines a device class.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
UPDATE DEVCLASS	Changes the attributes of a device class.

QUERY DOMAIN (Query a policy domain)

Use this command to display information on one or more policy domains.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query Domain-.*-----.-Format----Standard-----.
'-domain_name-' '-Format----+Standard+-'
'-Detailed-'
```

Parameters

domain_name

Specifies the policy domain to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy domains are displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display a summary of policy domains

Display partial information for all policy domains on the server. Issue the command:

```
query domain
```

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
EMPLOYEE- _RECORDS PROG1	VACATION	ACTIVEFI- LES	6	Employee Records Domain
PROG2			0	Programming Group Test Domain
STANDARD	STANDARD	STANDARD	1	Installed default policy domain

See Field descriptions for field descriptions.

Example: Display the list of active-data pools

Display the active-data pool list. Issue the command:

```
query domain format=detailed

    Policy Domain Name: STANDARD
    Activated Policy Set: STANDARD
    Activation Date/Time: 05/16/2006 16:18:05
    Days Since Activation: 15
    Activated Default Mgmt Class: STANDARD
    Number of Registered Nodes: 1
        Description: Installed default policy domain.
    Backup Retention (Grace Period): 30
    Archive Retention (Grace Period): 365
    Last Update by (administrator): SERVER_CONSOLE
        Last Update Date/Time: 05/31/2006 15:17:48
        Managing profile:
        Changes Pending: Yes
    Active Data Pool List: ADPPPOOL
```

See Field descriptions for field descriptions.

Field descriptions

Policy Domain Name

The name of the policy domain.

Activated Policy Set

The name of the policy set that was last activated in the domain.

The definitions in the last activated policy set and the ACTIVE policy set are not necessarily identical. When you activate a policy set, the server copies the contents of the policy set to the policy set with the special name ACTIVE. The copied definitions in the ACTIVE policy set can be modified only by activating another policy set. You can modify the original policy set without affecting the ACTIVE policy set. Therefore, definitions in the policy set that was last activated might not be the same as those in the ACTIVE policy set.

Activation Date/Time

The date and time that the policy set was activated.

Days Since Activation

The number of days since the policy set was activated.

Activated Default Mgmt Class

The assigned default management class for the policy set.

Number of Registered Nodes

The number of client nodes registered to the policy domain.

Description

The description of the policy domain.

Backup Retention (Grace Period)

The number of days to retain inactive backup versions of files when any of the following conditions occur:

- A file is rebound to a new management class, but neither the new management class nor default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.
- The backup copy group is deleted from the management class to which a file is bound and the default management class does not contain a backup copy group.

Archive Retention (Grace Period)

The number of days to retain an archive file that meets either of the following conditions:

- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.
- The archive copy group is deleted from the management class to which a file is bound and the default management class does not contain an archive copy group.

Last Update by (administrator)

Specifies that complete information is displayed for the drive.

Example: List drives associated with the server

Display information about all drives associated with the server. Issue the command:

```
query drive
```

Library Name	Drive Name	Device Type	Online
LIB1	DRIVE01	3590	Yes
LIB2	DRIVE02	3590	Yes

See Field descriptions for field descriptions.

Example: Display detailed information on a specific drive and library

Display detailed information about the drive named DRIVE02 that is associated with the library LIB2. Issue the command:

```
query drive lib2 drive02 format=detailed
```

```
Library Name: LIB2
Drive Name: DRIVE02
Device Type: 3590
On-Line: Yes
Drive State: Empty
Allocated to:
Last Update by (administrator): ADMIN
Last Update Date/Time: 02/29/2002 09:26:23
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE
```

See Field descriptions for field descriptions.

Field descriptions

Library Name

The name of the library to which the drive is assigned.

Drive Name

The name assigned to the drive.

Device Type

The device type as specified in the associated device class. The server must have a path defined from the server to the drive in order to determine the true device type. As long as there is a path defined from the server to the drive, the server will display the true device type of the drive even if there are other paths defined to this drive. Exceptions to this occur if the device type is remote or unknown.

REMOTE

The server does not have a path to the device. The only defined paths to the device are from data movers.

UNKNOWN

No path exists.

Tip: Review the output of the QUERY PATH command to determine if the desired paths are defined. If they are not defined, define those desired paths using the DEFINE PATH command. Also, if using a data mover device, review the output of the QUERY DATAMOVER command to determine the type of the data mover device. If you are using a path from the server to a drive, the device type of the device class and the drive need to match. If you are using a path from a data mover device to a drive, review the documentation for your type of data mover to ensure the device type of the device class is compatible with the type of data mover device.

On-Line

Specifies the status of the drive:

Yes

The drive is online and available for server operations.

No

The drive is offline and was put in this state by an administrator updating the status.

Unavailable Since

Specifies that the drive has been unavailable since *mm/dd/yy hh:mm:ss*. Output shows the time the server marked the drive as unavailable.

Polling Since

Specifies that the server is polling the drive because the drive stopped responding. Output shows the time the server detected a problem and began polling. The server polls a drive before stating it is unavailable. The time output follows the format: *mm/dd/yy hh:mm:ss*.

Read Formats

The read formats for the drive.

Write Formats

The write formats for the drive.

Element

The element number for the drive.

Drive State

This specifies the current state of this particular drive based on the result of the last SCSI command to the drive or library. The server tracks the state of the drive to improve its selection of a drive for an operation and its drive recovery operations. The values are:

Unavailable

The drive is not available to the library for operations.

Empty

The drive is empty and ready for operations.

Loaded

The drive is currently loaded, and the server is performing operations to the drive.

Unloaded

The media has been ejected from the drive.

Reserved

The drive is reserved for a mount request.

Unknown

The drive begins in drive state unknown as a result of being defined, as a result of server initialization, or as a result of having its status updated to online.

Volume Name

The volume name for the drive.

Allocated To

The name of the library client that is currently using the drive. This applies to shared SCSI libraries only; the field is left blank for all other libraries.

WWN

The World Wide Name for the drive.

Last Update by (administrator)

Who performed the last update to the drive.

Last Update Date/Time

The date and time when the last update occurred.

Cleaning Frequency (Gigabytes/ASNEEDED/NONE)

How often the server activates drive cleaning. This value can be the number of gigabytes, ASNEEDED, or NONE.

Related commands

Table 1. Commands related to QUERY DRIVE

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE DRIVE	Changes the attributes of a drive.

QUERY DRMEDIA (Query disaster recovery media)

Use this command to display information about database backup volumes, and volumes in copy storage pools, container-copy storage pools, and active-data storage pools. You can also use the command to create a file of executable commands to process the volumes.

The processing of volumes by this command depends on what the volumes are used for:

Backups of the server database

To control whether the command processes database backup volumes, use the SOURCE parameter. The command can process volumes that are used for full plus incremental or snapshot database backups. You cannot specify virtual volumes (backup objects that are stored on another server). You can change volumes through each state, or you can use the TOSTATE parameter and skip states to simplify the movements.

Copy storage pools

The QUERY DRMEDIA command always processes eligible copy storage-pool volumes.

Container-copy storage pools

By default, volumes in container-copy storage pools are not eligible for processing by the QUERY DRMEDIA command. To process container-copy storage pool volumes, you must issue the SET DRMCOPYCONTAINERSTGPOOL command first, or specify the COPYCONTAINERSTGPOOL parameter on the QUERY DRMEDIA command.

Active-data storage pools

By default, volumes in active-data storage pools are not eligible for processing by the QUERY DRMEDIA command. To process active-data pool volumes, you must issue the SET DRMACTIVEDATASTGPOOL command first, or specify the ACTIVEDATASTGPOOL parameter on the QUERY DRMEDIA command.

If you are using an external library and have moved a volume to the NOTMOUNTBLE state using the MOVE DRMEDIA command, the QUERY DRMEDIA command might still report the volume state as MOUNTABLE if it detects that the volume is in the library. Refer to the external library documentation for information about the procedures that you should follow when you use the MOVE DRMEDIA and the QUERY DRMEDIA commands.

Privilege class

To issue this command, you must have one of the following privilege classes:

- If the CMD parameter is NOT specified: operator or system privilege.
- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO: operator, unrestricted storage, or system privilege.
- If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES (the default): system privilege.

Syntax

```

      .-*-----
>>-Query DRMedia-+----->
                  '-volume_name-'

      .-WHEREState----All-----
>--+-----+----->
  '-WHEREState----+-----'
                +-All-----+
                +-MOUNTable-----+
                +-NOTMOUNTable----+
                +-COURier-----+
                +-VAULT-----+
                +-VAULTRetrieve----+
                +-COURIERRetrieve--+
                '-REmote-----'

>--+-----+-----+----->
  '-BEGINdate----date-' '-ENDDate----date-'

>--+-----+-----+----->
  '-BEGINTime----time-' '-ENDTime----time-'

>--+-----+-----+----->
  '-COPYstgpool----pool_name-'
```

```

>----->
'-ACTIVEDatastgpool---pool_name-'

>----->
'-COPYCONtainerstgpool---pool_name-'

.-Source---DBBackup-----.-Format---Standard-----.
>----->
'-Source---DBBackup---' '-Format---Standard---'
      +-DBSnapshot-+      +-Detailed-+
      '-DBNone-----'      '-Cmd-----'

>----->
'-WHERELOCation---location-' | .-----|
                              | V         | |
                              '-CMD-----"command"--+-'

                              .-APPend---No-----.
>-----><
'-CMDFilename---file_name-' '-APPend---No---'
                              '-Yes-'

```

Parameters

volume_name

Specifies the names of the volumes to be queried. You can use wildcard characters to specify multiple names. This parameter is optional. The server looks for matching names among the following eligible volumes:

- Database backup volumes, as selected by the SOURCE parameter of this command.
- Copy storage pool volumes from copy storage pools specified by the COPYSTGPOOL parameter. If you do not use the COPYSTGPOOL parameter, the server queries volumes from copy storage pools previously specified by the SET DRMCOPYSTGPOOL command.
- Active-data storage pool volumes from active-data storage pools specified by the ACTIVEDATASTGPOOL parameter. If you do not use the ACTIVEDATASTGPOOL parameter, the server queries volumes from active-data storage pools that were previously specified by the SET DRMACTIVEDATASTGPOOL command.
- Container-copy storage pool volumes from container-copy storage pools specified by the COPYCONTAINERSTGPOOL parameter. If you do not use the COPYCONTAINERSTGPOOL parameter, the server queries volumes from container-copy storage pools that were previously specified by the SET DRMCOPYCONTAINERSTGPOOL command.

Other parameters can also limit the results of the query.

WHEREState

Specifies the state of volumes to be processed. This parameter is optional. The default is ALL. Possible values are:

All

Specifies all volumes in all states.

Mountable

Volumes in this state contain valid data and are accessible for onsite processing.

NOTMountable

Volumes in this state are onsite, contain valid data, and not accessible for onsite processing.

COURier

Volumes in this state are being moved to an offsite location.

VAult

Volumes in this state are offsite, contain valid data, and are not accessible for onsite processing.

VAULTRetrieve

Volumes in this state are located at the offsite vault, do not contain valid data, and can be moved back onsite for reuse or disposal:

- A copy storage pool volume is considered to be in the VAULTRETRIEVE state if it has been empty for at least the number of days specified with the REUSEDELAY parameter on the DEFINE STGPOOL command.
- A database backup volume is considered to be in the VAULTRETRIEVE state if it is associated with a database backup series that was expired based on the value specified using the SET DRMDBBACKUPEXPIREDDAYS command.

Important: When you issue QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE, the server dynamically determines which volumes can be moved back onsite for reuse or disposal. Therefore, to ensure that you identify all volumes that are in a VAULTRETRIEVE state, issue QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE without the BEGINDATE,

ENDDATE, BEGINTIME or ENDTIME parameters. The *Last Update Date/Time* field in the output for QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE displays the date and time that a volume was moved to the VAULT state, not VAULTRETRIEVE.

COURIERRetrieve

Volumes in this state are being moved back to the onsite location.

REmote

Volumes in this state contain valid data and are located at the offsite remote server.

BEGINDate

Specifies the beginning date used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command has changed the volume to its current state on or after the specified date. The default is the earliest date for which volume information exists.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY-7 or -7. To query volumes beginning with records changed to their current state a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies the ending date used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command has changed the volume to its current state on or before the specified date. The default is the current date.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY-7 or -7. To query volumes beginning with records changed to their current state a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command has changed the volume to its current state on or after the specified time and date. The default is midnight (00:00:00) on the date specified with the BEGINDATE parameter.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	12:33:28
NOW	The current time on the specified begin date	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 <i>or</i> +03:00. If you issue QUERY DRMEDIA command at 9:00 with <code>BEGINTIME=NOW+03:00</code> or <code>BEGINTIME=+03:00</code> . The server displays volumes that were changed to their current state at 12:00 on the begin date that you specify.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-03:30 <i>or</i> -03:30. If you issue QUERY DRMEDIA command at 9:00 with <code>BEGINTIME=NOW-03:30</code> or <code>BEGINTIME=-03:30</code> . The server displays volumes that were changed to their current state at 5:30 on the begin date that you specify.

ENDTime

Specifies the ending time used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command has changed the volume to its current state on or before the specified time and date. The default is 23:59:59.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 <i>or</i> +03:00. If you issue QUERY DRMEDIA command at 9:00 with <code>ENDTIME=NOW+03:00</code> or <code>ENDTIME=+03:00</code> , IBM Spectrum Protect™ processes volumes that were changed to their current state at 12:00 on the end date you specify.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 <i>or</i> -03:30 If you issue QUERY DRMEDIA command at 9:00 with <code>ENDTIME=NOW-03:00</code> or <code>ENDTIME=-03:00</code> , IBM Spectrum Protect processes volumes that were changed to their current state at 6:00 on the end date you specify.

COPYstgpool

Specifies the name of the copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. The copy storage pools specified with this parameter override those specified with the SET DRMCOPYSTGPOOL command.

If this parameter is not specified, the server selects the storage pools as follows:

- If the SET DRMCOPYSTGPOOL command was previously issued with valid copy storage pool names, the server processes only those storage pools.
- If the SET DRMCOPYSTGPOOL command has not been issued, or if all of the copy storage pools have been removed using the SET DRMCOPYSTGPOOL command, the server processes all copy storage pool volumes in the specified state (ALL, MOUNTABLE, NOTMOUNTABLE, COURIER, VAULT, VAULTRETRIEVE, COURIERRETRIEVE, or REMOTE).

Source

Specifies whether any database backup volumes are selected. This parameter is optional. The default is DBBACKUP. Possible values are:

DBBackup

Full and incremental database backup volumes are selected.

DBSnapshot

Snapshot database backup volumes are selected.

DBNone

No database backup volumes are selected.

ACTIVEDatastgpool

Specifies the name of the active-data storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. The active-data storage pools that are specified with this parameter override those specified with the SET DRMACTIVEDATASTGPOOL command.

If this parameter is not specified, the server selects the storage pools as follows:

- If the SET DRMACTIVEDATASTGPOOL command was previously issued with valid active-data storage pool names, the server processes only those storage pools.
- If the SET DRMACTIVEDATASTGPOOL command has not been issued, or all of the active-data storage pools have been removed using the SET DRMACTIVEDATASTGPOOL command, the server processes all active-data storage pool volumes in the specified state (ALL, NOTMOUNTABLE, COURIER, VAULT, VAULTRETRIEVE, COURIERRETRIEVE, or REMOTE). Volumes in the MOUNTABLE state are not processed.

COPYCONtainerstgpool

Specifies the name of the container-copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. The container-copy storage pools that are specified using this parameter override those that are specified using the SET DRMCOPYCONTAINERSTGPOOL command.

If this parameter is not specified, the server selects the storage pools as follows:

- If the SET DRMCOPYCONTAINERSTGPOOL command was previously issued with names of valid container-copy storage pools, the server processes only those storage pools.
- If the SET DRMCOPYCONTAINERSTGPOOL command has not been issued, or if all container-copy storage pools were removed using the SET DRMCOPYCONTAINERSTGPOOL command, the server processes all container-copy pool volumes based on the value that is specified by the WHERESTATE parameter. If the parameter is set to a value of ALL, NOTMOUNTABLE, COURIER, VAULT, VAULTRETRIEVE, COURIERRETRIEVE, or REMOTE, the volumes are processed. If the value is set to MOUNTABLE, the volumes are not processed.

Format

Specifies the information to be displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that detailed information is displayed.

Cmd

Specifies that executable commands are built for the selected volumes. If you specify FORMAT=CMD, you must also specify the CMD parameter.

WHERELOcation

Specifies the location of the volumes to be queried. This parameter is optional. The maximum length of the location is 255 characters. Enclose the text in quotation marks if it contains any blank characters. If you specify a target server name, the disaster recovery manager displays all database backup volumes and copy storage pool volumes located on the target server.

CMD

Specifies the creation of executable commands to process the volume name and location obtained by this command. This parameter is optional. You must enclose the command specification in quotation marks. The maximum length of this parameter is 255 characters. The disaster recovery manager writes the commands to a file specified by the CMDFILENAME parameter or the SET DRMCMDFILENAME command, or generated by the QUERY DRMEDIA command. If the command length is greater than 240 characters, it is split into multiple lines and continuation characters (+) are added. You may need to alter the continuation character according to the product that runs the commands.

If you do not specify the FORMAT=CMD parameter, this command will not create any command lines.

string

The command string. The string must not include embedded quotation marks. For example, this is a valid CMD parameter:

```
cmd="checkin libvol lib8mm &vol status=scratch"
```

This is an example of a CMD parameter that is *not* valid:

```
cmd=""checkin libvolume lib8mm" &vol status=scratch""
```

substitution

Specifies a substitution variable to tell QUERY DRMEDIA to substitute a value for the variable. The variables are not case-sensitive, and must not contain blank spaces after the ampersand (&). The possible variables are:

&VOL

A volume name variable.

&LOC

A volume location.

&VOLDSN

The name of the file the server writes into the sequential access media labels. An example of a copy storage pool tape volume file name using the default prefix TSM is TSM.BFS. An example of a database backup tape volume file name using a prefix TSM310 defined with the device class is TSM310.DBB.

&NL

The new line character. When &NL is specified, QUERY DRMEDIA command splits the command at the &NL variable and does not append a continuation character. You must specify the proper continuation character before the &NL if one is required. If the &NL is not specified and the command line is greater than 240 characters, the line is split into multiple lines and continuation characters (+) are added.

AIX Linux CMDFilename

Specifies the fully qualified name of the file to contain the commands specified with CMD parameter. This parameter is optional.

If you do not specify a name with the SET DRMCMDFILENAME command, the server creates a file name by appending `exec.cmds` to the absolute directory path name of the IBM Spectrum Protect instance directory. If you specify a null string (""), the commands are displayed on the console only. You can redirect the commands to a file using the redirection character for the operating system.

If the operation fails after the command file is created, the file is not deleted.

Windows CMDFilename

Specifies the fully qualified name of the file to contain the commands specified with CMD parameter. This parameter is optional.

If you do not specify a file name with the SET DRMCMDFILENAME command, the server creates a file name by appending `exec.cmd` to the directory that represents this instance of the server (typically the directory where the IBM Spectrum Protect server was originally installed). If you specify a null string (""), the commands are displayed on the console only. You can redirect the commands to a file by using `>` or `>>` provided by the system. The disaster recovery manager allocates the file name specified or generated. If the file exists, the disaster recovery manager tries to use it and any existing data is overwritten.

If the operation fails after the command file is created, the file is not deleted.

APPend

Specifies whether to overwrite any existing contents of the command file or append the commands to the file. This parameter is optional. The default is NO. Possible values are:

No

The disaster recovery manager overwrites the contents of the file.

Yes

The disaster recovery manager appends the commands to the file.

Example: List volumes to be sent to offsite storage

Display all volumes to be given to a courier for offsite storage.

```
query drmedia wherestate=notmountable
format=standard
```

Volume Name	State	Last Update Date/Time	Automated LibName
-----	-----	-----	-----
TAPE01	Not mountable	01/20/1998 14:25:22	
DBTP01	Not mountable	01/20/1998 14:25:22	
DBTP03	Not mountable	01/20/1998 14:31:53	

See Field descriptions for field descriptions.

Example: Display information on volumes at the vault

Display detailed information about all volumes at the vault.

```
query drmedia wherestate=vault format=detailed

          Volume Name: DBTP02
                State: Vault
Last Update Date/Time: 01/20/1998 13:29:02
                Location: Ironmnt
                Volume Type: DBBackup
Copy Storage Pool Name:
Active-Data Storage Pool Name: TSMACTIVEPOOL
Automated LibName:
```

See Field descriptions for field descriptions.

Field descriptions

Volume Name

The name of the database backup or copy storage pool volume.

State

The state of the volume.

Last Update Date/Time

The date and time that the volume state was last updated. For volumes in the VAULTRETRIEVE state, this field displays the date and time that a volume was moved to the VAULT state, not VAULTRETRIEVE. The server does not "update" volumes to VAULTRETRIEVE. At the time the QUERY DRMEDIA command is issued, the server dynamically determines whether the data in copy storage pool volumes and database backup volumes is no longer valid and whether the volume can be brought back onsite for reuse or disposal.

Location

The Location field is displayed when the volume is not mountable or when it's not in the library. The Location field is empty if the volume is mountable and is in the library.

Volume Type

The type of volume. Possible values are:

DBBackup

A full or incremental database backup volume.

DBSnapshot

A database snapshot backup volume.

CopyStgPool

A copy storage pool volume.

ContcopyStgPool
A container-copy storage pool volume.

Copy Storage Pool Name
For a copy storage pool volume, the name of the copy storage pool.


Active Data Storage Pool Name
For an active-data storage pool volume, the name of the active-data storage pool.

Container-Copy Storage Pool Name
For a container-copy storage pool volume, the name of the container-copy storage pool.

Automated LibName
The name of the automated library if the volume is in a library.

Related commands

Table 1. Commands related to QUERY DRMEDIA

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
 SET DRMCOPYCONTAINERSTGPOOL	Specifies the container-copy storage pools that are used in DRM commands.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMDBBACKUPEXPIREDDAYS	Specifies criteria for database backup series expiration.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.

QUERY DRMSTATUS (Query disaster recovery manager system parameters)

Use this command to display information about the system parameters defined for disaster recovery manager (DRM).

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query DRMStatus-----><
```

Parameters

None.

Example: Display DRM system parameter information

Display information about the DRM system parameters:

query drmstatus

```
Recovery Plan Prefix:
Plan Instructions Prefix:
Replacement Volume Postfix: @
Primary Storage Pools: PRIM1 PRIM2
Copy Storage Pools: COPY*
Active-Data Storage Pools: TSMACTIVEPOOL
Container-Copy Storage Pools: COPYCNTRPOOL
Not Mountable Location name: Local
Courier Name: Fedex
Vault Site Name: Ironmnt
DB Backup Series expiration days: 30 Day(s)
Recovery Plan File Expiration Days: 30 Days(s)
Check Label?: No
Process FILE Device Type?: No
Command file name:
```

Field descriptions

Recovery Plan Prefix

User-specified prefix portion of the file name for the recovery plan file.

Plan Instructions Prefix

User-specified prefix portion of the file names for the server recovery instructions files.

Replacement Volume Postfix

The character added to the end of the replacement volume names in the recovery plan file.

Primary Storage Pools

The primary storage pools that are eligible for processing by the PREPARE command. If this field is blank, all primary storage pools are eligible.

Copy Storage Pools

The copy storage pools that are eligible for processing by the MOVE DRMEDIA, PREPARE, and QUERY DRMEDIA commands. If this field is blank, all copy storage pools are eligible.

Active-Data Storage Pools

The active-data pools that are eligible for processing by the MOVE DRMEDIA, PREPARE, and QUERY DRMEDIA commands. If this field is blank, active-data pools are not eligible.

Container-Copy Storage Pools

The container-copy storage pools that are eligible for processing by the MOVE DRMEDIA, PREPARE, and QUERY DRMEDIA commands. If this field is blank, container-copy storage pools are not eligible.

Not Mountable Location Name

The name of the offsite location where the media to be shipped are stored.

Courier Name

The name of the courier used to carry the media to the vault.

Vault Site Name

The name of the vault where the media is stored.

DB Backup Series Expiration Days

The minimum number of days that must elapse since a database series has been created before it is eligible to be expired. See the SET DRMDBBACKUPEXPIREDDAYS command for information about the criteria that must be met for database backup series expiration.

Recovery Plan File Expiration Days

The minimum number of days that must elapse since a recovery plan file, stored on a target server, has been created before it is eligible to be expired. See the SET DRMRPFEXPIREDDAYS command for information about the criteria that must be met for recovery plan file expiration.

Check Label?

Whether media labels are read for sequential media volumes checked out by the MOVE DRMEDIA command. Possible values are Yes or No.

Process FILE Device Type?

Whether MOVE DRMEDIA or QUERY DRMEDIA commands process database backup and copy storage pool volumes associated with a device class with a FILE device type. Possible values are Yes or No.

Command File Name

The full path file name that contains the executable commands generated by the MOVE DRMEDIA or QUERY DRMEDIA command.

Related commands

Table 1. Commands related to QUERY DRMSTATUS

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
SET DRMCHECKLABEL	Specifies whether IBM Spectrum Protect should read volume labels during MOVE DRMEDIA command processing.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
AIX Linux Windows SET DRMCOPYCONTAINERSTGPOOL	Specifies the container-copy storage pools that are used in DRM commands.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMCOURIERNAME	Specifies the name of the courier for the disaster recovery media.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.
SET DRMINSTRPREFIX	Specifies the prefix portion of the path name for the recovery plan instructions.
SET DRMPLANVPOSTFIX	Specifies the replacement volume names in the recovery plan file.
SET DRMPLANPREFIX	Specifies the prefix portion of the path name for the recovery plan.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SET DRMRPFEXPIREDAYS	Set criteria for recovery plan file expiration.
SET DRMVaultNAME	Specifies the name of the vault where DRM media is stored.
SET DRMNOTMOUNTABLENAME	Specifies the location name of the DRM media to be sent offsite.

QUERY ENABLED (Query enabled events)

Use this command to display either a list of enabled events or a list of disabled events, whichever is shorter.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query--ENabled--+-CONSOLE-----+----->
      +-ACTLOG-----+
      +-EVENTSERVER----+
      +-FILE-----+
      +-FILETEXT-----+
      |               (1) |
      +-NTEVENTLOG-----+
      |               (2) |
      +-SYSLOG-----+
      +-TIVOLI-----+
      '-USEREXIT-----'
```

```
>-----<
+-NODEname--==--node_name-----+
'-SERVername-----server_name-'
```

Notes:

1. This parameter is only available for the Windows operating system.
2. This parameter is only available for the Linux operating system.

Parameters

receiver

Specifies a type of receiver for enabled events. This is a required parameter. Valid values are:

ACTLOG

Specifies the IBM Spectrum Protect™ activity log as a receiver.

CONSOLE

Specifies the standard server console as a receiver.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

Windows NTEVENTLOG

Windows Specifies the Windows application log as a receiver.

Linux SYSLOG

Linux Specifies the Linux system log as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which IBM Spectrum Protect writes information as a receiver.

NODENAME

Specifies a node name to be queried. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for events enabled for the server running this command.

SERVername

Specifies a server name to be queried. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for events enabled for the server running this command.

Example: Query the server for console events

Query the server for server events that are enabled for the console. There are 10000 possible server events. Either a list of enabled events or disabled events is displayed (whichever list is shorter).

```
query enabled console
```

```
9998 events are enabled for the CONSOLE receiver. The
following events are DISABLED for the CONSOLE receiver:
```

```
ANR8409, ANR8410
```

Related commands

Table 1. Commands related to QUERY ENABLED

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.


```
'-EXceptiononly--==--No--+'
      '-Yes-'

.-Format----Standard----.
>--+-----+----->>
'-Format----+Standard--+'
      '-Detailed-'
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedules belong. You can use a wildcard character to specify this name.

schedule_name (Required)

Specifies the name of the schedule for which events are displayed. You can use a wildcard character to specify this name.

Type=Client

Specifies that the query displays events for client schedules. This parameter is optional. The default is CLIENT.

Nodes

Specifies the name of the client node that belongs to the specified policy domain for which events are displayed. You can specify multiple client nodes by separating the names with commas and no intervening spaces. You can use wildcard characters to specify nodes. If you do not specify a client name, events display for all clients that match the domain name and the schedule name.

BEGINDate

Specifies the beginning date of the time range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified	TODAY-7 or -7. To query events scheduled to start during the past seven days, specify BEGINDATE=TODAY-7 ENDDATE=TODAY or BEGINDATE=-7 ENDDATE=TODAY.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	10:30:08

Value	Description	Example
NOW	The current time on the specified begin date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either BEGINTIME=NOW+03:00 or BEGINTIME=+03:00. IBM Spectrum Protect™ displays events at 12:00 on the specified begin date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue this command at 9:00 to query events scheduled to start during the last 4 hours, you can specify either BEGINTIME=NOW-04:00 ENDTIME=NOW or BEGINTIME=-04:00 ENDTIME=NOW. IBM Spectrum Protect displays events at 5:00 on the specified begin date.

ENDDate

Specifies the ending date of the time range for events to be displayed. All events that were schedule to start during this time are displayed. This parameter is optional. The default is the value used for the BEGINDATE.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified	TODAY-8 or -8. To query events scheduled to start during a one-week period that ended yesterday, you can specify either BEGINDATE=TODAY-8 ENDDATE=TODAY-1 or BEGINDATE=-8 ENDDATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either <code>BEGINTIME=NOW ENDTIME=NOW+03:00</code> or <code>BEGINTIME=NOW ENDTIME=+03:00</code> .
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-04:00 or -04:00

EXceptiononly

Specifies the type of information you want on scheduled or completed events. This parameter is optional. The default is NO. You can specify one of the following values:

No

Specifies that the information on past and projected events is displayed.

Yes

Specifies that the events that failed or did not process as scheduled are displayed.

Format

Specifies how information displays. This parameter is optional. The default is STANDARD. The following values are possible:

Standard

Specifies that partial information for events displays.

Detailed

Specifies that complete information for events displays.

Display partial information for unsuccessful events

Display partial information for all events that are scheduled for DOMAIN1 that did not run successfully. Limit the search to the client named JOE. Limit the events that are displayed to events that were scheduled to occur from February 11, 2001 (02/11/2001) to February 12, 2001 (02/12/2001).

```
query event domain1 * nodes=joe begindate=02/11/2001
enddate=02/12/2001 exceptiononly=yes
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/11/1999 01:00:00	02/11/1999 01:13:55	BACK1	JOE	Failed
02/12/1999 01:00:00		DAILYBKP	JOE	Missed

See Field descriptions for field descriptions.

Display partial information for scheduled events for a client

Display complete information for all events that are scheduled for processing. Use the start time as 10 days previous to today, and the finish includes today.

```
query event * * begindate=today-10 enddate=today
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/04/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/04/2013 14:00:00	02/04/2013 14:12:49	VDATAMVR1-IN1	VDATAMVR1-T1	Completed
02/04/2013 14:30:00	02/04/2013 14:33:10	VDATAMVR1-IN2	VDATAMVR1-T2	Completed
02/04/2013 15:00:00	02/04/2013 15:01:49	VDATAMVR1-IN3	VDATAMVR1-T3	Completed
02/04/2013 15:30:00	02/04/2013 15:42:00	VDATAMVR1-IN4	VDATAMVR1-T4	Completed
02/05/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/05/2013 14:00:00	02/05/2013 14:05:22	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/05/2013 14:30:00	02/05/2013 14:32:53	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/05/2013 15:00:00	02/05/2013 15:00:38	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/05/2013 15:30:00	02/05/2013 15:36:41	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/06/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/06/2013 14:00:00	02/06/2013 14:06:42	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/06/2013 14:30:00	02/06/2013 14:35:41	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/06/2013 15:00:00	02/06/2013 15:08:56	VDATAMVR1-F3	VDATAMVR1-F3	Completed

02/06/2013 15:30:00	02/06/2013 15:40:49	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/07/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/07/2013 14:00:00	02/07/2013 14:03:43	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/07/2013 14:30:00	02/07/2013 14:35:10	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/07/2013 15:00:00	02/07/2013 15:09:12	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/07/2013 15:30:00	02/07/2013 15:40:21	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/08/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/08/2013 14:00:00	02/08/2013 14:10:17	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/08/2013 14:30:00	02/08/2013 14:39:16	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/08/2013 15:00:00	02/08/2013 15:08:17	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/08/2013 15:30:00	02/08/2013 15:41:16	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/09/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/09/2013 14:02:16		VDATAMVR1-F1	VDATAMVR1-F1	Failed 12
02/09/2013 14:30:00	02/09/2013 14:44:26	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/09/2013 15:00:00	02/09/2013 15:06:24	VDATAMVR1-F3	VDATAMVR1-F3	Failed 12
02/09/2013 15:30:00	02/09/2013 15:32:18	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/11/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/11/2013 14:00:00	02/11/2013 14:01:05	VDATAMVR1-F1	VDATAMVR1-F1	Failed 12
02/11/2013 14:30:00	02/11/2013 14:31:42	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/11/2013 15:00:00	02/11/2013 15:06:17	VDATAMVR1-F3	VDATAMVR1-F3	Failed 12
02/11/2013 15:30:00	02/11/2013 15:30:19	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/12/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/12/2013 14:00:00	02/12/2013 14:03:37	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/12/2013 14:30:00	02/12/2013 14:33:07	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/12/2013 15:00:00	02/12/2013 15:03:56	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/12/2013 15:30:00	02/12/2013 15:36:44	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/13/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/13/2013 14:00:00	02/13/2013 14:06:24	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/13/2013 14:30:00	02/13/2013 14:34:50	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/13/2013 15:00:00	02/13/2013 15:15:01	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/13/2013 15:30:00	02/13/2013 15:30:18	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/14/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Future
02/14/2013 14:00:00		VDATAMVR1-F1	VDATAMVR1-F1	Future
02/14/2013 14:30:00		VDATAMVR1-F2	VDATAMVR1-F2	Future
02/14/2013 15:00:00		VDATAMVR1-F3	VDATAMVR1-F3	Future

See Field descriptions for field descriptions.

Display detailed information for scheduled events for a client

Display the detailed information for events that are scheduled for processing by client DOC between the hours of 10:00 AM and 11:00 AM on November 1, 2005 (11/01/2005). Notice that when the status is FAILED, the result code is displayed.

```
query event domain1 * nodes=doc begindate=11/01/2005
begintime=10:00 endtime=11:00 enddate=11/01/2005
exceptionsonly=yes format=detailed
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
11/01/2005 10:01:01	11/01/2005 10:03:46	T1	DOC	Failed 8
11/01/2005 10:16:01	11/01/2005 10:16:10	T1	DOC	Failed 4
11/01/2005 10:31:01	11/01/2005 10:33:08	T1	DOC	Completed
11/01/2005 10:46:01		T1	DOC	Missed
11/01/2005 10:57:49	11/01/2005 10:58:07	T0	DOC	Failed 12

Field descriptions

Policy Domain Name

Specifies the name of the policy domain to which the schedule is assigned.

Schedule Name

Specifies the name of the schedule that initiated this event.

Node Name

Specifies the client that is scheduled to perform the operation.

Scheduled Start

Specifies the scheduled starting date and time for the event.

Actual Start

Specifies the date and time at which the client began processing the scheduled operation. No information is displayed if the scheduled operation has not started.

Completed

Specifies the date and time the scheduled event is completed.

Status

Specifies the status of the event at the time the QUERY EVENT command is issued. The following values are possible:

Completed

Specifies that the scheduled event is completed.

Failed

Specifies that the client reports a failure when you run the scheduled operation and successive retries failed.

Failed - no restart

Specifies an intermediate status, when a client session is interrupted by a communications error or timeout on the server. This status can be changed to a final status of "Completed" or "Failed" when the event completes.

Future

Specifies that the beginning of the startup window for the event is in the future. This status also indicates that an event record has not been created for this event.

In Progress

Specifies that the scheduled event is running and has not yet reported the completion state to the server.

Periodically check the status for completion of the scheduled event. If this status is not updated in a reasonable amount of time, review your client dsm Sched.log and dsmerror.log to determine why the client did not report the outcome of this event to the server. If the scheduled backup failed, rerun the scheduled event or perform a manual incremental backup to ensure the data backup.

Missed

Specifies that the scheduled startup window for this event passed and the schedule did not begin.

Pending

Specifies that the QUERY EVENT command was issued during the startup window for the event, but processing the scheduled operation did not begin.

Restarted

Specifies that the client has tried to process the scheduled operation again.

Severed

Specifies that the communications with the client is severed before the event can complete.

Started

Specifies that the event has begun processing.

Uncertain

Specifies that the state of the event cannot be determined. The server specifies `Uncertain` if the QUERY EVENT command does not find an event record. An event record is not found if the record was deleted or if the server was unavailable during the scheduled startup window (the schedule was never started). Records with `Uncertain` status are not stored in the database. If you do not want these records to display, either specify `EXCEPTIONSONLY=YES` or delete the schedule if it is no longer needed.

Attention: When a scheduled operation is processing, and is not restarted within its specified duration, the Status field shows `Started`. If the operation continues beyond the specified duration, no event record is created. If a query is issued after the specified duration has passed, the Status shows as `Failed` even if the operation is still running. After the operation completes, an event record is created, and a subsequent query shows the result in the Status field.

Result

Specifies the return code that indicates whether the schedule processed successfully. If the return code is a value other than 0, examine the server activity log and the client's error log and schedule log.

Return code	Explanation
0	All operations were completed successfully.
4	The operation was completed, but some files were not processed.
8	The operation was completed with at least one warning message.
12	The operation was completed with at least one error message. The count of error messages does not include notifications about skipped files.
-99	The operation failed because the session between the client and the server ended for an unknown reason. It is unknown whether the client can reconnect to the server to complete the schedule event.

If a schedule has `ACTION=COMMAND` as a parameter, and the command is not an IBM Spectrum Protect command, the command can produce other values in the Result field.

Reason

Specifies the reason for the return code.

QUERY EVENT (Display administrative event schedules)

Use the QUERY EVENT command to display scheduled and completed events for selected administrative command schedules.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query Evt--schedule_name--Type-----Administrative----->
.-BEGINdate-----current_date-. .-BEGINTime-----00:00-.
>--+-----+-----+-----+-----+-----+-----+----->
'-BEGINdate-----date-----' '-BEGINTime-----time--'

.-ENDDate-----begin_date-. .-ENDTime-----23:59-.
>--+-----+-----+-----+-----+-----+-----+----->
'-ENDDate-----date-----' '-ENDTime-----time--'

.-EXceptiononly-----No-----
>--+-----+-----+-----+-----+-----+-----+----->
'-EXceptiononly-----+No--+-'
                                     '-Yes-'

.-Format-----Standard-----
>--+-----+-----+-----+-----+-----+-----+-----><
'-Format-----+Standard+-'
                                     '-Detailed-'
```

Parameters

schedule_name (Required)

Specifies the name of the schedule for which events display. You can use wildcard characters to specify names.

Type=Administrative (Required)

Specifies that the query displays events for administrative command schedules.

BEGINDate

Specifies the beginning date of the time range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified	TODAY-7 or -7. To query events scheduled to start during the past seven days, specify BEGINDATE=TODAY-7 ENDDATE=TODAY or BEGINDATE=-7 ENDDATE=TODAY.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

Value	Description	Example
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either <code>BEGINTIME=NOW+03:00</code> or <code>BEGINTIME=+03:00</code> . IBM Spectrum Protect™ displays events at 12:00 on the specified begin date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue this command at 9:00 to query events scheduled to start during the last 4 hours, you can specify either <code>BEGINTIME=NOW-04:00</code> or <code>BEGINTIME=-04:00</code> . IBM Spectrum Protect displays events at 5:00 on the specified begin date.

ENDDate

Specifies the ending date of the time range for events to be displayed. All events that were schedule to start during this time are displayed. This parameter is optional. The default is the value used for the `BEGINDATE`.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or -days	The current date minus days specified	TODAY-8 or -8. To query events scheduled to start during a one-week period that ended yesterday, you can specify either <code>BEGINDATE=TODAY-8 ENDDATE=TODAY-1</code> or <code>BEGINDATE=-8 ENDDATE=-1</code> .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

Value	Description	Example
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either <code>BEGINTIME=NOW ENDTIME=NOW+03:00</code> or <code>BEGINTIME=NOW ENDTIME=+03:00</code> .
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-04:00 or -04:00

EXceptiononly

Specifies the type of information you want on scheduled or completed events. This parameter is optional. The default is NO. You can specify one of the following values:

No

Specifies that the information on past and projected events is displayed.

Yes

Specifies that the events that failed or did not process as scheduled are displayed.

Format

Specifies how the information displays. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information for events displays.

Detailed

Specifies that complete information for events displays.

Example: List events for a specific administrative schedule

Display partial information for all events scheduled for an administrative schedule named DOSADMIN. Limit the query to events that are scheduled for March 30, 1999 (03/30/1999). Issue the command:

```
query event dosadmin type=administrative
begindate=03/30/1999
enddate=03/30/1999
```

```
Scheduled Start      Actual Start      Schedule Status
-----
03/30/1999 00:00:00 03/30/1999 00:00:01  DOSADMIN  Completed
03/30/1999 04:00:00 03/30/1999 04:00:01  DOSADMIN  Completed
03/30/1999 12:00:00                DOSADMIN  Future
03/30/1999 16:00:00                DOSADMIN  Future
```

Field descriptions

Scheduled Start

Specifies the scheduled starting date and time for the event.

Actual Start

Specifies the date and time at which the client began processing the scheduled operation. No information displays if the schedule has not started executing.

Schedule Name

Specifies the name of the schedule that initiated this event.

Status

For administrative commands or scripts that specify WAIT=YES, the status of a scheduled event is STARTED until the operation specified by the command or script is completed. The final status of the scheduled event depends on the return code of the operation. However, if WAIT=YES and if the schedule is running a script that specifies PREVIEW=YES, the final status is COMPLETED, unless the script contained a syntax error.

For administrative commands or scripts that specify WAIT=NO, the status of a scheduled event is COMPLETED if the scheduled command or script started. The success of the schedule is independent of the success of the operation performed by the command or script.

QUERY EVENTRULES (Query rules for server or client events)

Use this command to display the history of events that are enabled or disabled by a specified receiver for the server or for a client node.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query--EVENTRules-----*----->>
| .-,-----|
| V          |
+---+-CONSOLE-----+
| ++ACTLOG-----+
| ++EVENTSERVER----+
| ++FILE-----+
| ++FILETEXT-----+
| | (1) |
| +-NTEVENTLOG-----+
| | (2) |
| ++SYSLOG-----+
| ++TIVOLI-----+
| '-USEREXIT-----'
+-NODEname-----node_name-----+
'-SERVername-----server_name-'
```

Notes:

1. This parameter is only available for the Windows operating system.
2. This parameter is only available for the Linux operating system.

Parameters

receivers

Specifies the name of one or more receivers for enabled events. This parameter is optional.

You can use a wildcard character to specify all receivers.

Valid values are:

CONSOLE

Specifies the standard console as a receiver.

ACTLOG

- Specifies the IBM Spectrum Protect™ activity log as a receiver.
- EVENTSERVER
 - Specifies the event server as a receiver.
- FILE
 - Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.
- FILETEXT
 - Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.
- Windows** NTEVENTLOG
 - Windows** Specifies the Windows application log as a receiver.
- Linux** SYSLOG
 - Linux** Specifies the Linux system log as a receiver.
- TIVOLI
 - Specifies the Tivoli Management Environment (TME) as a receiver.
- USEREXIT
 - Specifies a user-written routine to which IBM Spectrum Protect writes information as a receiver.
- NODENAME
 - Specifies a node name to be queried. You can use a wildcard character to specify a name. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for event rules for the server running this command.
- SERVER
 - Specifies a server name to be queried. You can use a wildcard character to specify a name. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for event rules for the server running this command.

Example: Display the history of client events for the server console

Display the history of client events enabled or disabled for the server console and activity log receivers.

```
query eventrules console,actlog nodename=*
```

Date/Time	Client Event Rules
05/29/97 13:39:58	ENABLE EVENTS CONSOLE ANE4001 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4962 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4963 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4965 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4966 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4967 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4968 NODENAMES=JEE
05/30/97 14:24:20	ENABLE EVENTS CONSOLE ANE4015 NODENAMES=RON
05/30/97 14:24:50	ENABLE EVENTS CONSOLE ANE4026 NODENAMES=DONNA
05/30/97 14:25:59	ENABLE EVENTS CONSOLE ANE4015 NODENAMES=DONNA

Example: Display the history of client events for all receivers

Display the history of server events enabled or disabled for all receivers.

```
query eventrules
```

Date/Time	Server Event Rules
05/22/97 14:35:13	ENABLE EVENTS CONSOLE ANR2578
05/30/97 14:29:31	ENABLE EVENTS CONSOLE ANR0272
05/30/97 14:31:46	ENABLE EVENTS USEREXIT ANR0130
05/30/97 14:31:54	ENABLE EVENTS USEREXIT ANR0131
05/30/97 14:50:28	ENABLE EVENTS USEREXIT ANR0266

Field descriptions

Date/Time

Specifies the date and time when the event was enabled or disabled.

Client Event Rules

Specifies client events that were enabled or disabled for the specified receivers.

Server Event Rules

Specifies server events that were enabled or disabled for the specified receivers.

Related commands

Table 1. Commands related to QUERY ENABLED

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.

QUERY EVENTSERVER (Query the event server)

Use this command to display the name of the event server.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query EVENTSErVer-----<<
```

Example: Display the event server name

Display the name of the event server.

```
query eventserver
```

```
ANR1669I Server EVENT is defined as the event server.
```

Related commands

Table 1. Commands related to QUERY EVENTSERVER

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DEFINE EVENTSERVER	Defines a server as an event server.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE EVENTSERVER	Deletes reference to the event server.
DELETE SERVER	Deletes the definition of a server.
END EVENTLOGGING	Ends event logging to a specified receiver.

QUERY EXPORT (Query for active or suspended export operations)

Use this command to list all restartable export operations. A restartable export is a server-to-server export operation whose FILEDATA value is not NONE. Only active server-to-server export operations that can be suspended are displayed.

Any EXPORT NODE or EXPORT SERVER operation with FILEDATA=NONE are not displayed. Additionally, the QUERY EXPORT command does not show export operations where the target device is either sequential media or virtual volumes.

Privilege class

An administrator can issue this command.

Syntax

```

>>-Query EXPort-----*----->
      +-----+
      |---export_identifier---|
      +-----+

.-State-----All-----
>--+-----+-----+----->
  |'-State-----+All-----+'
    +-Running---+
    '-SUSPended-'

>--+-----+-----+----->
  |'-PROcess-----process_number-'|
  +-----+

.-Format-----Standard-----
>--+-----+-----+----->>
  |'-Format-----+Standard-+'
    '-Detailed-'

```

Parameters

export_identifier

This optional parameter is the unique string identifier for the server-to-server export operation. Wildcard characters can be used to specify this name, and all matching export operations are queried. If you do not specify a value for this parameter and you also do not specify a PROCESS identifier, then all export operations are queried.

STate

This optional parameter queries the state of the valid server-to-server export operations. The default value is ALL. The possible values are:

ALL

Lists all running and suspended server-to-server export operations.

RUunning

Lists all active server-to-server export operations that are identifying eligible files or exporting files to the target server.

SUSPended

Lists all suspended server-to-server export operations. These suspended operations stopped running because of a failure or by the SUSPEND EXPORT command being issued.

PROcess

This optional parameter specifies the number of a running server-to-server export operation that you want to query. If PROCESS is specified, IBM Spectrum Protect™ only displays the running server-to-server export operation associated with the process number. If PROCESS is not specified, IBM Spectrum Protect displays information on all server-to-server export operations. You cannot specify this parameter if you specify an export identifier or if you specify the STATE parameter with a value of SUSPENDED.

Format

This optional parameter specifies how the information is displayed. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified export operations.

Detailed

When specified, displays all available information for the export operations.

Example: Display running and suspended export operations

List information for all currently running and suspended export operations. Issue the following command:

```
query export state=all
```

Export Identifier	Start Time	State	Process ID	Command
MYEXPORTNODE	01/24/2007 10:30:03	Suspended	--	Export NODE me,you,them filespace=c\$

```

nametype=unicode
filedata=all
durunits=indefinite
toserver=athens
exportid=MYEXPORTNODE

EXPORT_HOME_ 01/25/2007 Running 11 Export NODE n2,n3,n4
DIRS         09:30:03                filesystem=/home
                                         nametype=server
                                         filedata=all
                                         durunits=indefinite
                                         toserver=athens
                                         exportid=EXPORT_HOME_DIRS

EXPORT_NODE_ 01/25/2007 Running Not -- Export NODE n5,n6,n7
0001         14:30:33 Suspensible                filesystem=d$
                                         nametype=unicode
                                         filedata=archive
                                         durunits=indefinite
                                         toserver=athens

```

See Field descriptions for field descriptions.

Example: Display information about a running export operation

List information for the currently running export operation with process number "7." Issue the following command:

```

query export process=7

Export      Start Time  State   Process  Command
Identifier  -----
-----
MYEXPORTNODE 01/24/2007  Running  7        Export NODE
10:30:03                me,you,them
                        filesystem=c$
                        nametype=unicode
                        filedata=all
                        toserver=athens
                        exportid=MYEXPORTNODE

```

See Field descriptions for field descriptions.

Example: Display detailed information about all suspended export operations

List information for all currently suspended export operations. Issue the following command:

```

query export state=suspended format=detailed

Export Identifier: MyExportNode
Start Time: 01/24/2007 10:30:03
State: Suspended
Process Id: --
Command: Export NODE m* filesystem=c$
        nametype=unicode
        filedata=all durunits=indefinite
        toserver=athens
Phase: File list complete. Exporting
      eligible files
Total Running Time: 3 Days 0 Hours 24 Minutes
Current Process Running Time:
Export Operation Restart Count: 0
Date and Time of Last Restart: --
Date and Time of Last Suspend: 01/25/2007 08:30:11
Policy Domains Exported: 0
Policy Sets Exported: 0
Schedules Exported: 0
Mgmt Classes Exported: 0
Copy Groups Exported: 0
Administrators Exported: 1
Option Sets Exported: 0
Node Definitions Exported: 3
Filespace Definitions Exported: 7
Archive Files Exported: 50,000

```

```

Backup Files Exported: 150,000
Space Managed Files Exported: 0
Archive Files Skipped: 0
Backup Files Skipped: 25
Space Managed Files Skipped: 0
Total bytes Transferred (MB): 7,000
Total Files to be Transferred: 900,000
Files Remaining: 700,000

```

See Field descriptions for field descriptions.

Example: Display information for server-to-server export operations

List detailed information for all currently running server-to-server export operations. Issue the following command:

```

query export state=running format=detailed

Export Identifier: export_HOME_Dirs
Start Time: 01/25/2007 09:30:03
State: Running
Process Id: 11
Command: Export NODE n2,n3,n4
         filespace=/home nametype=
         server filedata=all
         toserver=athens
Phase: Identifying and exporting
       eligible files
Total Running Time: 0 Days 22 Hours 0 Minutes
Current Process Running Time: 01:30:00
Export Operation Restart Count: 4
Date and Time of last Restart: 02/01/2007 11:00:03
Date and Time of last Suspend: 01/31/2007 05:01:00
Policy Domains Exported: 0
Policy Sets Exported: 0
Schedules Exported: 0
Mgmt Classes Exported: 0
Copy Groups Exported: 0
Administrators Exported: 1
Option Sets Exported: 0
Node Definitions Exported: 3
Filespace Definitions Exported: 7
Archive Files Exported: 0
Backup Files Exported: 1000
Space Managed Files Exported: 0
Archive Files Skipped: 0
Backup Files Skipped: 0
Space Managed Files Skipped: 0
Total bytes Transferred (MB): 50
Total Files to be Transferred: 400,000
Files Remaining: 399,000

```

See Field descriptions for field descriptions.

Field descriptions

Export identifier

The unique identifier assigned to this server-to-server export operation.

Start time

The time and date that this export operation was first initiated.

State

The current state of this export operation. The value is one of the following:

Running - Not Suspending

The operation is active and is transmitting definitions to the target server. The process cannot be suspended, and if the process fails while in this state, you cannot restart it.

Running

The operation is active and is either searching for eligible files or transmitting file data to the target server.

Running - Suspend in Progress

The operation is in the process of being suspended as a result of a SUSPEND EXPORT command. The export operation is fully suspended when all of the data from the export operation is saved. An export operation in this state

does not respond to the following commands:

- CANCEL PROCESS
- CANCEL EXPORT
- RESTART EXPORT
- SUSPEND EXPORT

Suspended

The operation stopped running due to a failure or was suspended with the SUSPEND EXPORT command.

Process ID

The process ID for the export operation when the status is either "Initializing" or "Running".

Command

The full command issued to start this server-to-server export.

Phase

The current step that the operation is performing. The possible phases are shown in the order in which they are performed:

Creating definitions on target server

The operation is exporting definitions. The process cannot be suspended. Should the process fail in this phase, it cannot be restarted.

Identifying and exporting eligible files

The operation is building a list of eligible files for export. Some files may also be transmitted to the target during this phase. A process in this phase can be suspended. Should the process fail in this phase, it can be restarted.

File list complete. Exporting eligible files

The operation has completed building the list of eligible files for export and it is now transmitting the files to the target. A process in this phase can be suspended. Should the process fail in this phase, it can be restarted.

Total running time

The overall running time for this server-to-server export operation. For example, if this operation started and was then suspended and restarted two times, this value is the total running time of all three active processes of the export operation.

Current® process running time

The running time of the active process of a server-to-server export operation. No value is displayed for a suspended operation because no active process exists.

Export operation restart count

The number of times the server-to-server export operation was restarted.

Date and time of last restart

The last date and time at which this server-to-server export operation was restarted.

Date and time of last suspend

The last date and time at which this server-to-server export operation was suspended.

Policy domains exported

The number of policy domain definitions successfully exported to the target server.

Policy sets exported

The number of policy set definitions successfully exported to the target server.

Schedules exported

The number of schedule definitions successfully exported to the target server.

Mgmt classes exported

The number of management class definitions successfully exported to the target server.

Copy groups exported

The number of copy group definitions successfully exported to the target server.

Administrators exported

The number of administrator definitions successfully exported to the target server.

Option sets exported

The number of option set definitions successfully exported to the target server.

Node definitions exported

The number of node definitions successfully exported to the target server.

File space definitions exported

The number of file space definitions successfully exported to the target server.

Archive files exported

The number of archive files successfully exported to the target server.

Backup files exported

The number of backup files successfully exported to the target server.

Space managed files exported

The number of space managed files successfully exported to the target server.

Archive files skipped

The number of archive files that were eligible for export but were skipped.

Backup files skipped

The number of backup files that were eligible for export but were skipped.

Space managed files skipped

The number of space managed files that were eligible for export but were skipped.

Total bytes transferred (MB)

The total number of bytes transmitted so far to the target server for this export operation.

Total files to be transferred

The total number of files transmitted so far to the target server for this export operation.

Files remaining

The total number of files remaining to be transmitted to the target server for this export operation.

Related commands

Table 1. Commands related to QUERY EXPORT

Command	Description
CANCEL PROCESS	Cancels a background server process.
CANCEL EXPORT	Deletes a suspended export operation.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

AIX Linux Windows

QUERY EXTENTUPDATES (Query updated data extents)

Use this command to display information about updates to data extents in directory-container storage pools and to determine what data extents are deleted and what is eligible for deletion.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query EXTENTUPDates--pool_name-----<<
```

Parameters

pool_name (Required)

Specifies the storage pool to query. You cannot use wildcards to specify this name.

Example: Display information about updates to data extents

Display information about updates to data extents by issuing the following command:

```
query extentupdates
```


Parameters

node_name

Specifies the client node to which the file space belongs. You can use wildcard characters to specify this name. This parameter is optional. The default is all client node names.

You must specify a value for this parameter if you specify a file name.

file_space_name

Specifies the name of the file space to be queried. You can use wildcard characters to specify this name. This parameter is optional. If a value is not specified, all file spaces are queried.

If a server includes clients that use Unicode-enabled files spaces, the server might have to convert the name that you enter. For example, the server might have to convert the file space name that you enter from the server code page to Unicode. For more information, see the NAMETYPE parameter. If you do not specify a file space name, or if you specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

File space names are case-sensitive. You can use the QUERY FILESPACE command to determine the correct capitalization for the file space to be queried.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. You can specify one of the following values:

Standard

Specifies that partial information is displayed for the specified file space.

Detailed

Specifies that complete information is displayed for the specified file space.

NAMETYPE

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect™ clients that have Windows, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly qualified or fully qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret the file space names.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has problems accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODETYPE

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

Example: List all file spaces

Query all file spaces that are associated with all client nodes.

```
query filesystem
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity	Pct Util
JOE	\\joe\c\$	1	WinNT	NTFS	Yes	2,502.3	75.2
JOE	\\joe\d\$	2	WinNT	NTFS	Yes	6,173.4	59.6

See Field descriptions for field descriptions.

Example: Display detailed file space information for a virtual file space

Display detailed information for the file space /HomeDir, which is a virtual file space mapping and belongs to the NAS node NAS1.

```
query filesystem nas1 /HomeDir
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity	Pct Util
NAS1	/HomeDir	1	NetApp	WAFL (VFS)	No	2,502.3	75.2

See Field descriptions for field descriptions.

Important: You might not see the expected results after you request a detailed format because several fields must be completed by the API application. These fields include:

- File space type
- Platform
- Capacity
- Pct Util
- Last backup start Date/Time
- Last backup completion Date/Time

For more information about specific fields that are updated by the API, see the *IBM Spectrum Protect: Using the Application Programming Interface*.

Example: Display detailed file space information for a specific file space and node

Display detailed information about the \\joe\c\$ file space that belongs to the client node JOE.

```
query filesystem joe \\joe\c$ nametype=unicode format=detailed
```

```
Node Name: JOE
Filespace Name: \\joe\c$
Hexadecimal Filespace Name: 5c5c6a6f655c6324
FSID: 1
Collocation Group Name: FSGRP1
Platform: WinNT
Filespace Type: NTFS
Is Filespace Unicode?: Yes
Capacity: 2,502.3
Pct Util: 75.2
Last Backup Start Date/Time:
Days Since Last Backup Started:
Last Backup Completion Date/Time:
Days Since Last Backup Completed:
Last Replication Start Date/Time: 12/02/2012, 12:42:00
Days Since Last Node Replication Started: 30
Last Replication Completion Date/Time: 12/02/2012, 12:42:00
Days Since Last Replication Completed: 30
Last Backup Date/Time From Client (UTC): 06/02/2013, 09:10:00
Last Archive Date/Time From Client (UTC): 06/02/2013, 09:10:00
Backup Replication Rule Name: ACTIVE_DATA
Backup Replication Rule State: ENABLED
Archive Replication Rule Name: DEFAULT
Archive Replication Rule State: ENABLED
Space Management Replication Rule Name: NONE
Space Management Replication Rule State: DISABLED
```


At-risk type: Custom interval
At-risk interval: 2,222
Decommissioned: No
Decommissioned Date:
MAC Address:

See Field descriptions for field descriptions.

Field descriptions

Important: You might not see the expected results after requesting a detailed format because several fields must be completed by the API application. These fields include:

- Filespace Type
- Platform
- Capacity
- Pct Util
- Last Backup Start Date/Time
- Last Backup Completion Date/Time

For more information about specific fields that are updated by the API, see the *IBM Spectrum Protect: Using the Application Programming Interface*.

Node Name

Specifies the name of the client node.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Hexadecimal Filespace Name

Specifies the hexadecimal name of the file space for the client node in UTF-8 format.

FSID

Specifies the file space ID of the file space.

Collocation Group Name

The name of the collocation group, if any, to which the file space belongs.

Platform

Specifies the platform for the client node.

Filespace Type

Specifies the type of file space.

A file space type that is appended with "(VFS)" denotes that this file space name is a virtual file space mapping for a directory path on a NAS device.

Is Filespace Unicode?

Indicates whether the file space is Unicode.

Capacity

Specifies the amount of space, in megabytes, assigned to this file space on the client node.

For a file space that is a virtual file space mapping for a directory path, this field represents the capacity of the file space on which the directory path is located.

Pct Util

Specifies the percentage of the file space that is occupied.

For a file space that is a virtual file space mapping for a directory path, the percentage used is calculated as the percentage of the capacity of the file space that was occupied by the directory at the time of the last full backup.

Last Backup Start Date/Time

Specifies the start date and time of the last incremental backup of the file space.

Days Since Last Backup Started

Specifies the number of days since the start of the last incremental backup of the file space.

Last Backup Completion Date/Time

Specifies the completion date and time of the last incremental backup of the file space.

Days Since Last Backup Completed

Specifies the number of days since the completion of the last incremental backup of the file space.

Last Replication Start Date/Time

Specifies the date and time that the last replication of file space data started.

Days Since Last Replication Started

Specifies the number of days since the last replication of file space data started.

Last Replication Completion Date/Time

Specifies the date and time that the last replication of file space data ended.

Days Since Last Replication Completed

Specifies the number of days since the last replication of file space data ended.

Last Backup Date/Time From Client (UTC)

The date and time, in Universal Time Coordinates (UTC), of the last backup operation for this file space.

Last Archive Date/Time From Client (UTC)

The date and time, in Universal Time Coordinates (UTC), of the last archive operation for this file space.

Backup Replication Rule Name

Specifies the replication rule that applies to backup data in the file space. The following values are possible:

ALL_DATA

Replicates active and inactive backup data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority.

Attention: If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the FORCERECONCILE=YES parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data. The data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority.

DEFAULT

Replicates backup data according to the client node rule for backup data. If the client node rule for backup data is DEFAULT, backup data is replicated according to the server rule for backup data.

NONE

Backup data in the file space is not replicated.

Backup Replication Rule State

Specifies whether replication of backup data in the file space is enabled or disabled. If the state is ENABLED, backup files are eligible for replication. If the state is DISABLED, backup files are not eligible for replication.

Archive Replication Rule Name

Specifies the replication rule that applies to archive data in the file space. The following values are possible:

ALL_DATA

Replicates archive data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates archive data. The data is replicated with a high priority.

DEFAULT

Replicates archive data according to the client rule for archive data. If the client rule for archive data is DEFAULT, archive data is replicated according to the server rule for archive data.

NONE

Archive data in the file space is not replicated.

Archive Replication Rule State

Specifies whether replication of archive data in the file space is enabled or disabled. If the state is **ENABLED**, archive files are eligible for replication. If the state is **DISABLED**, archive files are not eligible for replication.

Space Management Replication Rule Name

Specifies the replication rule that applies to space-managed data in the file space. The following values are possible:

ALL_DATA

Replicates space-managed data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates space-managed data. The data is replicated with a high priority.

DEFAULT

Replicates space-managed data according to the client rule for space-managed data. If the client rule for space-managed data is **DEFAULT**, space-managed data is replicated according to the server rule for space-managed data.

NONE

Space-managed data in the file space is not replicated.

Space Management Replication Rule State

Specifies whether replication of space-managed data in the file space is enabled or disabled. If the state is **ENABLED**, space-managed files are eligible for replication. If the state is **DISABLED**, space-managed files are not eligible for replication.

At-risk type

Specifies the at-risk evaluation type. Values can be **Default**, **Bypassed**, or **Custom**. **Default** indicates that the node is evaluated with the same interval that was specified for the nodes classification by the **SET STATUSATRISKINTERVAL** command. **Bypassed** indicates that the node is not evaluated for at-risk status by the status monitor. **Custom** indicates that the node is evaluated with the interval that was specified by the **SET VMATRISKINTERVAL** command, rather than the interval that was specified by the **SET STATUSATRISKINTERVAL** command.

At-risk interval

Specifies the amount of time, in hours, between client backup activity before the status monitor considers the client at-risk. This field applies only when the at-risk type is **Custom**.

Decommissioned

Specifies whether the virtual machine that the file space represents is decommissioned.

Decommissioned Date

Specifies the date that the virtual machine that the file space represents was decommissioned.

MAC Address

Specifies the media access control (MAC) address of the file spaces backed up for VMWare virtual machines. In the case where the virtual machine has multiple MAC addresses this is the lowest valued address.

Related commands

Table 1. Commands related to QUERY FILESPACE

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
RENAME FILESPACE	Renames a client filesystem on the server.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes that are associated with a client node.

QUERY FSCOUNTS (Query number of objects)

Use this command to display information about the number of objects (files and directories) in file spaces that belong to a client node.

Tip: To obtain accurate information, issue the QUERY FSCOUNTS command after the backup operations ends. Also, if you are currently expiring objects from the file space, the numbers might not reflect the latest changes.

The database is queried and the counts are completed in real time.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*-----
      |               .-*----- |
>>-Query FSCounts--+node_name--+----->
                        '-file_space_name-'

.-Format---Standard----.  .-NAMEType---SERVER-----.
>--+-----+-----+-----+-----+----->
  '-Format---+Standard-+-'  '-NAMEType---+SERVER---+-'
                        '-Detailed-'                +-UNICODE-+
                                                '-FSID----'

.-CODEType---BOTH-----
>--+-----+-----+-----+-----+-----><
  '-CODEType---+UNICODE---+-'
                        +-NONUNICODE-+
                        '-BOTH-----'
```

Parameters

node_name

Specifies the client node to which the file space belongs. You can use wildcard characters to specify this name, or use a group name. A group name specifies the name of the group to which the client node belongs. This parameter is required. Comma-delimited lists are not allowed. An asterisk specifies all client nodes.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect™ clients that have Windows, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly qualified or fully qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret the file space names.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has problems accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specifies what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

Field descriptions

Node Name

Specifies the name of the client node.

FSID

Specifies the file space ID of the file space.

Filespace Type

Specifies the type of file space.

A file space type that is appended with "(VFS)" denotes that this file space name is a virtual file space mapping for a directory path on a network-attached storage (NAS) device.

Is Filespace Unicode?

Indicates whether the file space is Unicode.

Related commands

Table 1. Commands related to QUERY FSCOUNTS

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY OCCUPANCY	Displays file space information by storage pool.

QUERY LIBRARY (Query a library)

Use this command to display information about libraries.

Privilege class

Any administrator can issue this command.

Syntax

```
.*-----.  
>>-Query LIBRARY--+----->  
    '-library_name-'  
  
.-Format----Standard----.  
>--+-----+----->>  
    '-Format----+Standard-+-'  
        '-Detailed-'
```

Parameters

library_name

Specifies the name of the library to be queried. You can use wildcards to specify names. This parameter is optional.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the library.

Detailed

Specifies that complete information is displayed for the library.

Example: Display summary information about a specific library

Display information about the library named AUTO. Issue the command:

```
query library auto
```

```
Library Name: AUTO
Library Type: SCSI
  ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
  Shared: No
  LanFree:
ObeyMountRetention:
```

See Field descriptions for field descriptions.

Example: Display detailed library information about a specific library

Display information in full detail about the library named EZLIFE. Issue the command:

AIX

Linux

```
query library ezlife format=detailed
```

AIX

Linux

```
Library Name: EZLIFE
Library Type: SCSI
  ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
  Shared: Yes
  LanFree:
ObeyMountRetention:
Primary Library Manager: EZSERVER
  WWN:
  Serial Number:
  AutoLabel: OVERWRITE
Relabel Scratch: Yes
Last Update by (administrator): DOCTOR_MIKE
Last Update Date/Time: 2002-12-05 15:24:53
```

Windows

```
Library Name: EZLIFE
Library Type: SCSI
  ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
  Shared: YES
  LanFree:
ObeyMountRetention:
Primary Library Manager: EZSERVER
  WWN:
  Serial Number:
  AutoLabel: OVERWRITE
Reset Drives: No
Relabel Scratch: Yes
Last Update by (administrator): DOCTOR_MIKE
Last Update Date/Time: 2000-12-05 15:24:53
```

See Field descriptions for field descriptions.

Field descriptions

Library Name

The name of the library.

Library Type

The type of library.

ACS Id

Specifies that the library is a StorageTek library that is controlled by StorageTek Automated Cartridge System Library Software (ACSL).

Private Category

The category number for private volumes that must be mounted by name.

The information that is displayed in this field applies only to an IBM® 3494 or 3495 Tape Library Dataserver.

Scratch Category

The category number to use for scratch volumes in the library.

The information that is displayed in this field applies only to an IBM 3494 or 3495 Tape Library Dataserver.

WORM Scratch Category

The category number that is used for WORM scratch volumes in the library.

The information that is displayed in this field applies only to an IBM 3494 or 3495 Tape Library Dataserver.

External Manager

The location of the external library manager where the server can send media access requests.

Shared

Whether this library is shared with other IBM Spectrum Protect™ servers in a storage area network (SAN).

LanFree

Whether an external library is used for LAN-free operations.

ObeyMountRetention

Whether the server uses the value that is set for mount retention in the device class that is associated with this external library.

Primary Library Manager

The name of the server that is responsible for controlling access to library resources.

WWN

The Fibre Channel worldwide name for the library.



Serial Number

Specifies the serial number for the library that is being queried.

AutoLabel

Specifies whether the server attempts to automatically label tape volumes.

Reset Drives

  Specifies whether the server completes a target reset when the server is restarted or when a library client or storage agent re-connection is established.

Relabel Scratch

Specifies whether the server relabels volumes that were deleted and returned to scratch.

Last Update by (administrator)

Who completed the last update to the library.

Last Update Date/Time

The date and time when the last update occurred.

Related commands

Table 1. Commands related to QUERY LIBRARY

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE LIBRARY	Deletes a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE LIBRARY	Changes the attributes of a library.

QUERY LIBVOLUME (Query a library volume)

Use this command to display information about one or more volumes that are checked into an automated library for use by the IBM Spectrum Protect™ server.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*----- .-*-----
>>-Query LIBVolume-----+-----+----->
      '-library_name-' '-volume_name-'

      .-Format----Standard----.
>--+-----+-----+----->>
      '-Format----+Standard--+'
              '-Detailed-'

```

Parameters

library_name

Specifies the name of the library. You can use wildcard characters to specify this name. This parameter is optional. The default is all libraries.

volume_name

Specifies the volume name. You can use wildcard characters to specify this name. This parameter is optional. The default is all volumes.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List checked in volumes for a specific library

Display information about all of the volumes that are checked into the library named TAPE. See Field descriptions for field descriptions.

```
query libvolume tape
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element	Device Type
TAPE	000114	Scratch			1,000	LTO
TAPE	NY1602	Scratch			1,001	DLT

Example: Display detailed information for a specific library

Display detailed information about a volume named JJY008. See Field descriptions for field descriptions.

```
query libvolume jjy008 format=detailed
```

```

      Library Name: HPW3494
      Volume Name: JJY008
      Status: Private
      Owner: SUNSET
      Last Use: Data
      Home Element:
      Device Type:
Cleanings Left:
      Media Type:

```

Field descriptions

Library Name

The name of the library where the storage volume is located.

Volume Name

The name of the storage volume.

Status

The status of the storage volume according to the library inventory. If the status is Private, the volume is being used by IBM Spectrum Protect. If the status is Scratch, the volume is available for use.

Owner

The owner server of the volume, if the volume is private.

Last Use

The type of data on the volume. This field applies only to volumes in Private status. For storage pool volumes, this field shows **Data**. For database backup volumes (full, incremental, or snapshot), this field shows **DbBackup**.

Home Element

The element address of the library slot containing the volume.

Device Type

The type of device that the volume is used on. This field will display a value only for volumes checked into a library that has mixed media capabilities.

Cleanings Left

For cleaner cartridges, the number of cleanings left.

Media Type

The type of media the volume represents (for example, 8mm tape).

Related commands

Table 1. Commands related to QUERY LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE LIBVOLUME	Changes the status of a storage volume.

QUERY LICENSE (Display license information)

Use this command to display license audit, license terms, and compliance information.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query LICense-----><
```

Parameters

None.

To display the license information, issue the following command:

```
query license
```

The following example output is displayed:

```

ANR2017I Administrator
SERVER_CONSOLE issued command: QUERY LICENSE
Last License Audit: 10/17/2016
14:28:08
Number of Data Protection for Oracle in use: 0
Number of Data Protection for
Oracle in try buy mode: 0
Number of Data Protection for Microsoft SQL in use: 0
Number of Data Protection for
Microsoft SQL in try buy mode: 0
Number of Data Protection for
Microsoft Exchange in use: 0
Number of Data Protection for
MS Exchange in try buy mode: 0
Number of TDP for Lotus Notes in use: 12
Number of TDP for Lotus Notes in try buy mode: 0
Number of Data Protection for Lotus Domino in use: 0
Number of Data Protection for
Lotus Domino in try buy mode: 0
Number of TDP for Informix in use: 1
Number of TDP for Informix in try buy mode: 0
Number of TDP for SAP R/3 in use: 0
Number of TDP for SAP R/3 in try buy mode: 0
Number of TDP for ESS in use: 0
Number of TDP for ESS in try buy mode: 0
Number of TDP for ESS R/3 in use: 0
Number of TDP for ESS R/3 in try buy mode: 0
Number of TDP for EMC Symmetrix in use: 0
Number of TDP for EMC Symmetrix in try buy mode: 0
Number of TDP for EMC Symmetrix R/3 in use: 6
Number of TDP for EMC Symmetrix R/3 in try buy mode: 0
Number of TDP for WAS in use: 0
Number of TDP for WAS in try buy mode: 0
Is IBM Spectrum Protect for Data Retention in use?: No
Is IBM Spectrum Protect for Data Retention licensed?: Yes
Is IBM Spectrum Protect Basic Edition in use: Yes
Is IBM Spectrum Protect Basic Edition licensed: Yes
Is IBM Spectrum Protect Extended Edition in use: No
Is IBM Spectrum Protect Extended Edition licensed: Yes
Server License Compliance: Valid

```

Field descriptions

Last License Audit

Specifies the date and time when the last license audit occurred.

Number of Data Protection for Oracle in use

Specifies the number of Data Protection for Oracle that are in use. A product is in use if you purchased the product and registered the license.

Number of Data Protection for Oracle in try buy mode

Specifies the number of Data Protection for Oracle that are in try buy mode.

Number of Data Protection for Microsoft SQL in use

Specifies the number of Data Protection for Microsoft SQL that are in use. A product is in use if you purchased the product and registered the license.

Number of Data Protection for Microsoft SQL in try buy mode

Specifies the number of Data Protection for Microsoft SQL that are in try buy mode.

Number of Data Protection for Microsoft Exchange in use

Specifies the number of Data Protection for Microsoft Exchange that are in use. A product is in use if you purchased the product and registered the license.

Number of Data Protection for Microsoft Exchange in try buy mode

Specifies the number of Data Protection for Microsoft Exchange that are in try buy mode.

Number of TDP for Lotus Notes® in use

Specifies the number of TDP for Lotus Notes that are in use. A product is in use if you purchased the product and registered the license.

Number of TDP for Lotus Notes in try buy mode

Specifies the number of TDP for Lotus Notes that are in try buy mode.

Number of Data Protection for Lotus® Domino® in use

Specifies the number of Data Protection for Lotus Domino that are in use. A product is in use if you purchased the product and registered the license.

- Number of Data Protection for Lotus Domino in try buy mode
Specifies the number of Data Protection for Lotus Domino that are in try buy mode.
- Number of TDP for Informix® in use
Specifies the number of TDP for Informix that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for Informix in try buy mode
Specifies the number of TDP for Informix that are in try buy mode.
- Number of TDP for SAP R/3 in use
Specifies the number of TDP for SAP R/3 that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for SAP R/3 in try buy mode
Specifies the number of TDP for SAP R/3 that are in try buy mode.
- Number of TDP for ESS in use
Specifies the number of TDP for ESS that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for ESS in try buy mode
Specifies the number of TDP for ESS that are in try buy mode.
- Number of TDP for ESS R/3 in use
Specifies the number of TDP for ESS R/3 that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for ESS R/3 in try buy mode
Specifies the number of TDP for ESS R/3 that are in try buy mode.
- Number of TDP for EMC Symmetrix in use
Specifies the number of TDP for EMC Symmetrix that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for EMC Symmetrix in try buy mode
Specifies the number of TDP for EMC Symmetrix that are in try buy mode.
- Number of TDP for EMC Symmetrix R/3 in use
Specifies the number of TDP for EMC Symmetrix R/3 that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for EMC Symmetrix R/3 in try buy mode
Specifies the number of TDP for EMC Symmetrix R/3 that are in try buy mode.
- Number of TDP for WAS in use
Specifies the number of TDP for WAS that are in use. A product is in use if you purchased the product and registered the license.
- Number of TDP for WAS in try buy mode
Specifies the number of TDP for WAS that are in try buy mode.
- Is IBM Spectrum Protect™ for Data Retention in use ?
Specifies whether the IBM Spectrum Protect for Data Retention is in use. A product is in use if you purchased the product and registered the license.
- Is IBM Spectrum Protect for Data Retention licensed ?
Specifies whether the IBM Spectrum Protect for Data Retention is licensed.
- Is IBM Spectrum Protect Basic Edition in use
Specifies whether the IBM Spectrum Protect Basic Edition is in use. A product is in use if you purchased the product and registered the license.
- Is IBM Spectrum Protect Basic Edition licensed
Specifies whether the IBM Spectrum Protect Basic Edition is licensed.
- Is IBM Spectrum Protect Extended Edition in use
Specifies whether the IBM Spectrum Protect Extended Edition is in use. A product is in use if you purchased the product and registered the license.
- Is IBM Spectrum Protect Extended Edition licensed
Specifies whether the IBM Spectrum Protect Extended Edition is licensed.
- Server License Compliance
Specifies whether the server license is valid.

Related commands

Table 1. Commands related to QUERY LICENSE

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.

Command	Description
QUERY AUDITOCUPANCY	Displays the server storage utilization for a client node.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PVUESTIMATE	Displays processor value unit estimates. Remember: The QUERY PVUESTIMATE command reports licenses by providing PVU information on a per-node basis for server devices.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the IBM Spectrum Protect server.
REGISTER NODE	Defines a client node to the server and sets options for that user.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.
UPDATE NODE	Changes the attributes that are associated with a client node.

QUERY LOG (Display information about the recovery log)

Use this command to display information about the recovery log.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query LOG-----Format-----Standard----->>
'-Format-----+--Standard--+'
'-Detailed-'

```

Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. The following values are possible:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary information about the recovery log

Display summary information about the recovery log. See Field descriptions for field descriptions.

```
query log
```

```

Total Space (MB)    Used Space (MB)    Free Space (MB)
-----
38,912              543.3              38,368.7

```

AIX | Linux

Example: Display detailed information about the recovery log

Display detailed information about the recovery log. See Field descriptions for field descriptions.

```
query log format=detailed

    Active Log Directory : /actlog
      Total Space (MB): 524,032
      Used Space (MB): 3,517
      Free Space (MB): 520,515

Total Size of File System (MB): 564,443
Used Space on File System (MB): 527,049
Free Space on File System (MB): 8,722

    Archive Log Directory : /archlog
Total Size of File System (MB): 603,751.82
Used Space on File System (MB): 80,642.30
Free Space on File System (MB): 523,109.52
  Archive Log Compressed : Yes

    Mirror Log Directory : /mirrorlog
Total Size of File System (MB): 564,443
Used Space on File System (MB): 527,049
Free Space on File System (MB): 8,722

Archive Failover Log Directory : /archfaillog
Total Size of File System (MB): 301,372.06
Used Space on File System (MB): 44,741.80
Free Space on File System (MB): 256,630.26
```

Windows

Example: Display detailed information about the recovery log when the mirror log and the archive failover log are not defined

The output of this command on Windows systems is different. For example, the output contains blanks for the mirror log and the archive failover log.

Display information about the recovery log when the mirror log and the archive failover log are not defined.

```
query log format=detailed

    Active Log Directory : d:\actlog
      Total Space (MB): 524,032
      Used Space (MB): 3,517
      Free Space (MB): 520,515

Total Size of File System (MB): 564,443
Used Space on File System (MB): 527,049
Free Space on File System (MB): 8,722

    Archive Log Directory : e:\archlog
Total Size of File System (MB): 603,751.82
Used Space on File System (MB): 80,642.30
Free Space on File System (MB): 523,109.52
  Archive Log Compressed: Yes

    Mirror Log Directory :
Total Size of File System (MB):
Used Space on File System (MB):
Free Space on File System (MB):

Archive Failover Log Directory :
Total Size of File System (MB):
Used Space on File System (MB):
Free Space on File System (MB):
```

Windows

Field descriptions

- Total Space
Specifies the maximum size of the active log, in megabytes.
- Used Space
Specifies the amount of used active log space, in megabytes.
- Free Space
Specifies the amount of active log space that is not being used by uncommitted transactions, in megabytes.
- Total Size of File System
Specifies the total size of the file system, in megabytes.
- Space Used on File System
Specifies the amount of used space on the file system, in megabytes.
- Free Space on File System
Specifies the amount of space that is available on the file system, in megabytes.
- Archive Log Compressed
Specifies whether the archive logs are compressed.
- Active Log Directory
Specifies the location where active log files are stored. When you change the active log directory, the server moves all archived logs to the archive log directory and all active logs to a new active log directory.
- Mirror Log Directory
Specifies the location where the mirror for the active log is maintained.
- Archive Failover Log Directory
Specifies the location into which the server saves archive logs if the logs cannot be archived to the archive log directory.
- Archive Log Directory
Specifies the location into which the server can archive a log file after all the transactions that are represented in that log file are completed.

QUERY MACHINE (Query machine information)

Use this command to display information for one or more machines. You can use this information to recover IBM Spectrum Protect™ client machines in case of a disaster.

Attention: IBM Spectrum Protect does not use the information in any way. It is available only to help you plan for the disaster recovery of client machines.

IBM Spectrum Protect displays information for multiple machines in the following order:

- According to the priority specified.
- Within a priority, according to the specified location and machine name.

Privilege class

Any administrator can issue this command.

Syntax

```

.*-----
>>-Query MACHine-----+-----+----->
      '-machine_name-'   '-BUilding---building-'

>+-----+-----+----->
      '-FLoor---floor-'   '-ROom---room-'

>+-----+-----+----->
      '-PRIority---priority-'   '-ADSMServer---+Yes-+-'
                                   '-No--'

.-Format---Standard-----
>+-----+-----+----->>
      '-Format---+Standard-----+
              +-Detailed-----+
              +-RECOVERYInstructions-+
              '-CHaracteristics-----'
```

Parameters

machine_name

Specifies the name of one or more machines to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is all machines that meet the specified criteria.

BUilding

Specifies the name or number of the building that the machines are in. This parameter is optional. Enclose the text in quotation marks if it contains any blank characters.

FLOor

Specifies the name or number of the floor that the machines are on. This parameter is optional. Enclose the text in quotation marks if it contains any blank characters.

ROom

Specifies the name or number of the room that the machines are in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRiority

Specifies the priority number of the machines. This parameter is optional.

ADSMServer

Specifies if the machine contains an IBM Spectrum Protect server. This parameter is optional. The default is to display any machines that meet the other criteria. Possible values are:

Yes

The machine contains an IBM Spectrum Protect server.

No

The machines do not contain an IBM Spectrum Protect server.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Displays partial information for the machines.

Detailed

Displays all information for the machines.

RECOVERYInstructions

Displays only machine recovery instructions. This option is valid only when querying a specific machine.

CHaracteristics

Displays only machine characteristics. This option is valid only when querying a specific machine.

Example: Display information for a specific machine

Display information for a machine named MACH1. See Field descriptions for field descriptions.

```
query machine MACH1
```

Machine Name	Machine Priority	Building	Floor	Room	Node Name	Recovery Media Name
MACH1	1	21	2	2929	VIRGINIA	RECMED1

Example: Display detailed information for priority 1 machines

Display detailed information for all priority 1 machines on the second floor of building 21. See Field descriptions for field descriptions.

```
query machine * building=21 floor=2 priority=1  
format=detailed
```

```
Machine Name: MACH1  
Machine Priority: 1  
Building: 21  
Floor: 2  
Room: 2929  
Server?: Yes  
Description: TSM server machine  
Node Name: VIRGINIA  
Recovery Media Name: RECMED1  
Characteristics?: Yes  
Recovery Instructions?: Yes
```

Field descriptions

Machine Name	The name of the machine.
Machine Priority	The recovery priority of the machine.
Building	The building in which the machine is located.
Floor	The floor on which the machine is located.
Room	The room in which the machine is located.
Server?	Whether the machine contains an IBM Spectrum Protect server.
Description	A description of the machine.
Node Name	The IBM Spectrum Protect client nodes associated with this machine.
Recovery Media Name	The recovery media associated with this machine.
Characteristics?	Whether the characteristics text of the machine is stored in the database.
Recovery Instructions?	Specifies whether recovery instructions text for a machine is stored in the IBM Spectrum Protect database.

Related commands

Table 1. Commands related to QUERY MACHINE

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DEFINE MACHNODEASSOCIATION	Associates an IBM Spectrum Protect node with a machine.
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE MACHINE	Deletes a machine.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the IBM Spectrum Protect database.
UPDATE MACHINE	Changes the information for a machine.

QUERY MEDIA (Query sequential-access storage pool media)

Use this command to display information about the sequential-access primary and copy storage pool volumes moved by the MOVE MEDIA command.

Privilege class

Any administrator with system or operator privilege can issue this command unless it includes the CMD parameter. If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, unrestricted storage, or system privilege. If the CMD parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES (the default), the administrator must have system privilege.

The QUERY MEDIA command displays only volumes with an ACCESS MODE value of READONLY or READWRITE.

Syntax

```
>>-Query MEDIA-+-----+---STGpool-----pool_name----->
                .*-----
                '-volume_name-'
```


READOnly

Specifies that volumes with an access mode of READONLY display.

Format

Specifies how information displays. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information displays for the specified sequential access storage pool volumes.

Detailed

Specifies that complete information displays for the specified sequential access storage pool volumes.

Cmd

Specifies that executable commands are built for the storage pool volumes processed by the QUERY MEDIA command. These commands will be in the file specified with the CMDFILENAME parameter on the QUERY MEDIA command. If you want the commands to display on the console only, specify a null string ("" for the CMDFILENAME. If FORMAT=CMD is specified but no command string is specified with the CMD parameter, the QUERY MEDIA command will fail.

WHEREState

Specifies the state of volumes to process. This parameter restricts processing to volumes that have the specified state. This parameter is optional. The default is ALL. Possible values are:

All

Specifies that volumes in all states are queried. The valid states are: MOUNTABLEINLIB and MOUNTABLENOTINLIB.

MOUNTABLEInlib

Specifies that volumes that are currently in the MOUNTABLEINLIB state are queried. Volumes in the MOUNTABLEINLIB state are in the library, and are onsite, contain valid data, and are available for onsite processing.

MOUNTABLENotinlib

Specifies that volumes that are currently in the MOUNTABLENOTINLIB state are queried. Volumes in the MOUNTABLENOTINLIB state are not in the library, do not contain valid data, and are not available for onsite processing.

WHEREOVFLocation

Specifies the overflow location of the volumes to display. This parameter is optional. This parameter restricts processing to volumes that are in the specified location. The maximum length of the location is 255 characters. The location must be enclosed in quotation marks if it contains any blank characters.

CMd

Specifies the creation of executable commands. Enclose the command specification in quotation marks. The maximum length of the command specification is 255 characters. This parameter is optional.

For each volume successfully processed by the QUERY MEDIA command, the server writes the associated commands to a file. Specify the file name with the CMDFILENAME parameter.

AIX Linux If you do not specify a filename, the command will generate a default filename by appending the string `exec.cmds.media` to the server directory.

Windows If you do not specify a filename, the command will generate a default filename by appending the string `exec.cmd.media` to the server directory.

Remember:

1. If the command written to the file exceeds 255 characters, it is split into multiple lines, and a continuation character (+) is added to all but the last line. You may need to alter the continuation character according to the requirements of the product that runs the commands.
2. If an executable command is specified with any value for FORMAT other than CMD, the command string is ignored, and the QUERY MEDIA command will not write any command line.

Specify a command string and any substitution variables:

string

Specifies the string to build an executable command to process the volume name or volume location or both. You can specify any free form text for the string. Do not use embedded quotation marks. For example, the following is a valid executable command specification:

```
cmd="checkin libvolume &vol"
```

The following is an invalid executable command specification:

```
cmd="checkin libvolume "&vol""
```

substitution

Specifies a variable for which you want the QUERY MEDIA command to substitute a value. The possible substitution variables are:

&VOL

Substitute the volume name for &VOL. You can specify lowercase characters, &vol. No spaces or blanks are allowed between ampersand, &, and VOL. If there are spaces or blanks between ampersand and VOL, the QUERY MEDIA command will treat them as strings and no substitution will be set. If &VOL is not specified, no volume name is set in the executable command.

&LOC

Substitute the volume location for &LOC. You can specify lowercase characters, &loc. No spaces or blanks are allowed between ampersand, &, and LOC. If there are spaces or blanks between ampersand and LOC, the QUERY MEDIA command will treat them as strings and no substitution will be set. If &LOC is not specified, no location name is set in the executable command.

&VOLDSN

Substitute the volume file name for &VOLDSN. An example of a copy storage pool tape volume file name using the defined prefix IBM Spectrum Protect™ 310 is IBM Spectrum Protect310.BFS. If &VOLDSN is not specified, no volume file name is set in the executable command.

&NL

Substitute the new line character for &NL. When &NL is specified, the QUERY MEDIA command will split the command at the position where the &NL is and will not append any continuation character. The user is responsible for specifying the proper continuation character before the &NL if one is required. The user is also responsible for the length of the line written. If the &NL is not specified and the command exceeds 255 characters, the command is split into multiple lines, and a continuation character (+) is added to all but the last line.

CMDFilename

Specifies the full path name that will contain the commands specified with CMD parameter when FORMAT=CMD is specified. This parameter is optional. The maximum length of the file name is 1279 characters.

AIX | **Linux** If you specify "" with the CMDFILENAME parameter, the QUERY MEDIA command will generate a file name by appending the "exec.cmds.media" to the server directory. The server directory is the current working directory of the server process.

Windows If you specify "" with the CMDFILENAME parameter, the QUERY MEDIA command will generate a file name by appending the "exec.cmd.media" to the server directory. The server directory is the current working directory of the server process.

If you specify a null string ("") for the CMDFILENAME, the commands built are displayed on the console only. You can redirect the commands displayed to a file by using the redirection characters for the operating system (> or >>).

AIX | **Linux** If the filename is not specified, the command will generate a default filename by appending the string "exec.cmds.media" to the server directory.

Windows If the filename is not specified, the command will generate a default filename by appending the string "exec.cmd.media" to the server directory.

The QUERY MEDIA command automatically allocates the file name specified or generated. If the file name exists, the QUERY MEDIA command will attempt to use it and the existing data, if any, in the file to be overwritten. You can specify APPEND=YES to prevent the existing data from being overwritten. If the QUERY MEDIA command fails after the command file is allocated, the file is not deleted.

APPend

Specifies to write at the beginning or the ending of the command file data. This parameter is optional. The default is NO. Possible values are:

No

Specifies to write the data from the beginning of the command file. If the given command file exists, its contents are overwritten.

Yes

Specifies to append the command file by writing at the end of the command file data.

Example: Display information on a specific sequential access storage pool

Display all full and partial full volumes that are in the sequential access primary storage pool, ARCHIVE. See Field descriptions for field descriptions.

```
query media * stgpool=archive wherestatus=full, filling
```

Volume Name	State	Location	Automated LibName
TAPE01	Mountable in Library		LIB3494
TAPE03	Mountable not in Lib.	Room1234/Bldg31	
TAPE07	Mountable in Library		LIB3494
TAPE09	Mountable not in Lib.	Room1234/Bldg31	

Example: Display information on sequential access storage pool with a specific prefix

Display in detail all full volumes in MOUNTABLENOTINLIB state for sequential access storage pools that have a prefix name of ONSITE. See Field descriptions for field descriptions.

```
query media wherestate=mountablenotinlib stgpool=onsite*
wherestatus=full format=detailed
```

```
Volume Name: TAPE21
State: Mountable not in library
Volume Status: Full
Access: ReadOnly
Last Reference Date: 01/30/98
Last Update Date/Time: 08/20/1996 13:29:02
Location: Rm569/bldg31
Storage Pool Name: ONSITE.ARCHIVE
Automated Libname:
```

```
Volume Name: TAPE22
State: Mountable not in library
Volume Status: Full
Access: ReadOnly
Last Reference Date: 01/30/98
Last Update Date/Time: 08/20/1996 15:29:02
Location: Rm569/bldg31
Storage Pool Name: ONSITE.ARCHIVEPOOL
Automated Libname:
```

Example: Generate checkin commands

Generate the CHECKIN LIBVOLUME commands for full and partially full volumes that are in the ONSITE.ARCHIVE primary storage pool and stored in the overflow location Room 2948/Bldg31.

```
query media * stgpool=onsite.archive format=cmd
wherestatus=full,filling wherestate=mountablenotinlib
whereovflocation=room2948/bldg31
cmd="checkin libvol lib3494 &vol status=private"
cmdfilename=/tsm/move/media/checkin.vols
```

The QUERY MEDIA command created the CHECKIN LIBVOLUME executable commands in /tsm/move/media/checkin.vols, which can be run by issuing the MACRO command with /tsm/move/media/checkin.vols as the macro name.

```
checkin libvol lib3494 TAPE04 status=private
checkin libvol lib3494 TAPE13 status=private
checkin libvol lib3494 TAPE14 status=private
```

Field descriptions

Volume Name

Specifies the name of the primary sequential access storage pool volume.

State

Specifies the state of the volume.

Volume Status

Specifies the status of the volume.

Access

Specifies the access mode of the volume.

Last Reference Date

Specifies the volume's last written date or last read date, whichever is more recent.

Last Update Date/Time

Specifies the date and time when the volume was most recently updated.

Location

Specifies where the volume is stored. If the volume is ejected from the library and its location is not specified or defined, a question mark (?) is displayed for the location.

Storage Pool Name

Specifies the name of the sequential access storage pool where the volume is defined.

Automated LibName

Specifies the automated library name if the volume is in the library.

Related commands

Table 1. Commands related to QUERY MEDIA

Command			Description
AIX	Linux	Windows	MOVE MEDIA
			Moves storage pool volumes that are managed by an automated library.

QUERY MGMTCLASS (Query a management class)

Use this command to display information about management classes.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query Mgmtclass----->
.---*---*-----
>--+-----+----->
|          .---*---*-----|
|'-domain_name-----+'
|          .---*---*-----|
|          '-policy_set_name-----+'
|          '-class_name-'
.---Format---Standard-----
>--+-----+----->>
|'-Format---Standard-----+'
|          '-Detailed-'
```

Parameters

domain_name

Specifies the policy domain associated with the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, management classes in all policy domains are queried. You must specify this parameter when querying an explicitly named management class.

policy_set_name

Specifies the policy set associated with the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, management classes in all policy sets are queried. You must specify this parameter when querying an explicitly named management class.

class_name

Specifies the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all management classes are queried.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display information for all management classes

Query all management classes for all policy domains. Create the output in standard format. See Field descriptions for field descriptions.

```
query mgmtclass
```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Default Mgmt Class ?	Description
EMPLOYEE-RECORDS	ACTIVE	ACTIVEFILES	Yes	Modified default management class
EMPLOYEE-RECORDS	HOLIDAY	ACTIVEFILES	Yes	Modified default management class
EMPLOYEE-RECORDS	HOLIDAY	FILEHISTORY	No	Test modified management class
EMPLOYEE-RECORDS	VACATION	ACTIVEFILES	Yes	Original default management class
EMPLOYEE-RECORDS	VACATION	FILEHISTORY	No	Test modified management class
PROG1	SUMMER	MCLASS1	No	Technical Support Mgmt Class
PROG2	SUMMER	MCLASS1	No	Technical Support Mgmt Class
STANDARD	ACTIVE	STANDARD	Yes	Installed default management class
STANDARD	STANDARD	STANDARD	Yes	Installed default management class

To display information about management classes in a specific policy domain, for example the domain ENGPOLDOM, issue the following command:

```
query mgmtclass engpoldom * *
```

Example: Display detailed information for a specific management class

Query the ACTIVEFILES management class that is assigned to the VACATION policy set of the EMPLOYEE_RECORDS policy domain. Create the output in detailed format. See Field descriptions for field descriptions.

```
query mgmtclass employee_records vacation  
activefiles format=detailed
```

```
Policy Domain Name: EMPLOYEE_RECORDS  
Policy Set Name: VACATION  
Mgmt Class Name: ACTIVEFILES  
Default Mgmt Class ?: Yes  
Description: Installed default management class  
Space Management Technique: None  
Auto-Migrate on Non-Use: 0  
Migration Requires Backup?: Yes  
Migration Destination: SPACEMGPOOL  
Last Update by (administrator): $$CONFIG_MANAGER$$  
Last Update Date/Time: 05/31/1998 13:15:45  
Managing Profile: EMPLOYEE  
Changes Pending: Yes
```

Field descriptions

Policy Domain Name

The policy domain.

Policy Set Name

The policy set.

- Mgmt Class Name**
The management class.
- Default Mgmt Class ?**
Whether the management class is the default management class for the policy set.
- Description**
The description of the management class.
- Space Management Technique**
The space management technique for the management class, for IBM Spectrum Protect™ for Space Management clients.
- Auto-Migrate on Non-Use**
The number of days that must elapse since a file was last accessed before it is eligible for automatic migration by IBM Spectrum Protect for Space Management clients.
- Migration Requires Backup?**
Whether a backup version of a file must exist before a file can be migrated by IBM Spectrum Protect for Space Management clients.
- Migration Destination**
The storage pool that is the destination for files migrated by IBM Spectrum Protect for Space Management clients.
- Last Update by (administrator)**
The administrator or server that most recently updated the management class. If this field contains \$\$CONFIG_MANAGER\$\$, the management class is associated with a domain that is managed by the configuration manager.
- Last Update Date/Time**
The date and time when the management class was most recently defined or updated.
- Managing profile**
The profile or profiles to which the managed server subscribed to get the definition of this management class.
- Changes Pending**
Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 1. Commands related to QUERY MGMTCLASS

Command	Description
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY DOMAIN	Displays information about policy domains.
UPDATE MGMTCLASS	Changes the attributes of a management class.

QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)

Use this command to display information about alert monitoring and server status settings.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query MONITORSEttings-----<<
```

Display monitoring settings

Display details about the monitoring settings. See Field descriptions for more details.

query monitorsettings

Example output:

```
Monitor Status: On
Status Refresh Interval (Minutes): 5
Status Retention (Hours): 48
Monitor Message Alerts: On
Alert Update Interval (Minutes): 10
Alert to Email: On
Send Alert Summary to Administrators: On
Alert from Email Address: DJADMIN@MYDOMAIN.COM
Alert SMTP Host: DJHOST.MYDOMAIN.COM
Alert SMTP Port: 25
Alert Active Duration (Minutes): 480
Alert Inactive Duration (Minutes): 480
Alert Closed Duration (Minutes): 60
Monitoring Admin: ADMIN
Monitored Group: MONGROUP
Monitored Servers: SERVER2
At-Risk Interval for Applications: 24
Skipped files as At-Risk for Applications?: Yes
At-Risk Interval for Virtual Machines: 24
Skipped files as At-Risk for Virtual Machines?: Yes
At-Risk Interval for Systems: 24
Skipped files as At-Risk for Systems?: Yes
Deployment Repository: /source/packages/deploy
Maximum Deployment Packages: 4
Deployment Package Manager: On
```

Field descriptions

Monitor Status

Specifies whether alert monitoring on the server is enabled or disabled.

Status Refresh Interval (Minutes)

Specifies the number of minutes between intervals that the monitoring server gathers event data.

Status Retention (Hours)

Specifies the number of hours that status monitoring indicators are retained.

Monitor Message Alerts

Specifies whether alerts are sent to administrators by email.

Alert Update Interval (Minutes)

Specifies the length of time, in minutes, that the alert monitor waits before the alert is updated and pruned on the server.

Alert to Email

Specifies whether alerts are sent to administrators by email.

Send Alert Summary to Administrators

Specifies the administrators that receive a summary of existing alerts on the server in an email.

Alert from Email Address

Specifies the email address of the sender.

Alert SMTP Host

Specifies the Simple Mail Transfer Protocol (SMTP) host mail server that is used to send alerts by email.

Alert SMTP Port

Specifies the SMTP mail server port that is used to send alerts by email.

Alert Active Duration (Minutes)

Specifies how long, in minutes, an alert remains active.

Alert Inactive Duration (Minutes)

Specifies how long, in minutes, an alert remains inactive.

Alert Closed Duration (Minutes)

Specifies how long, in minutes, an alert remains closed before it is deleted from the server.

Monitoring Admin

Specifies the name of the monitoring administrator that is used to connect to the servers in the monitored group.

Monitored Group

Specifies the name of the monitored server group.

Monitored Servers

Specifies the names of the servers in the monitored server group. The monitor settings might be different on each monitored server. If so, issue the query command for each server to display the monitoring settings.

At-Risk Interval for Applications

Specifies how long, in hours, an applications client can log no activity before it is considered at-risk.
 Skipped files as At-Risk for Applications?
 Specifies that the server considers skipped files, by the client as a failure, and marks the client at-risk.
 At-Risk Interval for Virtual Machines
 Specifies how long, in hours, a virtual client can log no activity before it is considered at-risk.
 Skipped files as At-Risk for Virtual Machines?
 Specifies that the server considers skipped files, by the client as a failure and marks the client at-risk.
 At-Risk Interval for Systems
 Specifies how long, in hours, a systems client can log no activity before it is considered at-risk.
 Skipped files as At-Risk for Systems?
 Specifies that the server considers skipped files, by the client as a failure, and marks the client at-risk.
 Deployment Repository
 Specifies the location where client deployment packages are downloaded, and the location of the storage volumes that are used for client deployment packages.
 Maximum Deployment Packages
 Specifies the maximum number of client deployment packages that are stored in the deployment repository for each product version.
 Deployment Package Manager
 Specifies whether the deployment package manager queries the FTP site for new deployment packages and downloads new packages as they become available.

Related commands

Table 1. Commands related to QUERY MONITORSETTINGS

Command	Description
DEFINE ALERTTRIGGER (Define an alert trigger)	Associates specified messages to an alert trigger.
DELETE ALERTTRIGGER (Remove a message from an alert trigger)	Removes a message number that can trigger an alert.
DELETE GRPMEMBER (Delete a server from a server group)	Deletes a server from a server group.
DELETE SERVER (Delete a server definition)	Deletes the definition of a server.
QUERY ALERTSTATUS (Query the status of an alert)	Displays information about alerts that have been issued on the server.
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	Displays message numbers that trigger an alert.
SET ALERTMONITOR (Set the alert monitor to on or off)	Specifies whether alert monitoring is set to on or off.
SET DEPLOYREPOSITORY (Set the download path for client deployment packages)	Specifies the location where client deployment packages are downloaded.
SET DEPLOYMAXPKGS (Set the maximum number of client deployment packages to store)	Specifies the maximum number of client deployment packages that are downloaded and stored on the server.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE ALERTTRIGGER (Update a defined alert trigger)	Updates the attributes of one or more alert triggers.
UPDATE ALERTSTATUS (Update the status of an alert)	Updates the status of a reported alert.

QUERY MONITORSTATUS (Query the monitoring status)

Use this command to display monitoring messages that are within the defined status retention period.

You can limit the output to a specified status, such as only messages with a status of active. If you do not specify any parameters, all messages are displayed.

Privilege class

Any administrator can issue this command.

Syntax

```

.-Format---Standard-----
>>-Query MONITORStatus----->
    '-Format---Standard+-'
        '-Detailed-'

.-Type---Active-----
>----->
    '-Type---All-----' '-ACTivity---activity_name-'
        ++Active---+
        '-Inactive-'

>-----><
    '-Name---element_name-' | .,----- |
        |                   V           | |
        '-Status---Normal---+-'
            ++Warning+
            '-Error---'

```

Parameters

Format

Specifies the amount of information that is displayed. The default value is STANDARD. Specify one of the following values:

Standard

Specifies that only partial information is displayed for the specified messages.

Detailed

Specifies that all information is displayed for the specified messages.

Type

This parameter restricts the output to only messages with the specified type value. Specify one of the following values:

ALL

Displays all information.

Active

Displays all active messages. This is the default value.

Inactive

Displays all inactive messages.

ACTivity

Specifies the activity that you want to query. See the DEFINE STATUSTHRESHOLD command for details on available activities to query.

NAme

Specifies the name that you want to query. The NAME value refers to the name of the element with the specified activity. For example, a status indicator that contains information about a storage pool that is called `backuppool` has the NAME set to BACKUPPOOL.

STatus

Specifies the status of the messages that you want to query. You can specify multiple status values in a list by separating the values with commas and no intervening spaces. If you do not specify a value for this parameter, information for all status values is displayed. Specify one of the following values:

Normal

Displays all messages with a normal status.

Warning

Displays all messages with a warning status.

Error

Displays all messages with an error status.

Display monitoring settings

Display details about the monitoring status.

Query MONITORStatus type=active

Example output:

```
Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
```

Display monitoring settings

Display details about the monitoring status.

query monitorstatus f=d type=active

Example output:

```
Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY DISK AND
FILE STORAGE
Element Name: CAPACITY OF PRIMARY DISK AND
FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY DISK AND
FILE STORAGE
Element Name: USED CAPACITY OF PRIMARY DISK AND
FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
```

Second Alternate Repair Suggestion:

```
Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:
```

```
Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY
TAPE STORAGE
Element Name: USED CAPACITY OF PRIMARY
TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:
```

Field descriptions

Server Name

The name of the server.

Activity Date

The last date and time activity was reported.

Activity Name

The name of the activity.

Element Name

The name of the element.

Element Numeric Value

The numeric value of the element.

Element String Value

The string value of the element.

Element State

The state of the element.

Element Details

The detailed information of the element.

Primary Repair Suggestion

The primary repair suggestion.

First Alternate Repair Suggestion

The repair suggestion to follow if the primary suggestion is not adequate.

Second Alternate Repair Suggestion

The repair suggestion to follow if the primary and first alternate suggestions are not adequate.

Related commands

Table 1. Commands related to QUERY MONITORSTATUS

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.

Command	Description
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

QUERY MOUNT (Display information on mounted sequential access volumes)

Use this command to display information about the status of one or more sequential access volumes that are mounted.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query Mount-----.-Format----Standard----->>
      '-volume_name-' '-Format-----Standard--'
                                '-Detailed-'

```

Parameters

volume_name

Specifies the name of the mounted sequential access volume. You can use wildcard characters to specify this name. This parameter is optional. The default is all mounted volumes.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List all mounted sequential volumes

Display information on all mounted sequential media volumes.

```
query mount
```

AIX

```

ANR8330I 3590 volume D6W992 is mounted R/O
in drive RMT1(/dev/rmt1), status: IN USE.
ANR8334I 1 volumes found.
ANR8331I 8MMTAPE volume WPD000 is mounted R/W
in drive 8MM.1 (/dev/mt0), status: DISMOUNTING.
ANR8334I 1 volumes found.

```

Linux

```
ANR8330I 3590 volume D6W992 is mounted R/O
in drive RMT1/dev/IBMtape1, status: IN USE.
ANR8334I 1 volumes found.
ANR8331I 8MMTAPE volume WPD000 is mounted R/W
in drive 8MM.1 (/dev/tsm SCSI/mt0), status: DISMOUNTING.
ANR8334I 1 volumes found.
```

Windows

```
ANR8330I 3590 volume D6W992 is mounted R/O
in drive RMT1(/dev/rmt1), status: IN USE.
ANR8334I 1 volumes found.
ANR8331I 8MMTAPE volume WPD000 is mounted R/W
in drive 8MM.1 (mt3.0.0.0), status: DISMOUNTING.
ANR8334I 1 volumes found.
```

Remember:

1. If the status of a volume is full or if its access mode is read-only (R/O), the mount mode of the volume is R/O. To determine the status and access mode of a volume, issue the `QUERY VOLUME FORMAT=DETAILED` command. If a volume can be written to (that is, the status is filling or empty), the mount mode of the volume is read/write (R/W), even if it is only being read.
2. In a storage pool that is associated with the FILE or CENTERA device type, the server can complete concurrent multiple read-access and one write-access to the same volume. As a result, a volume in a storage pool with a device type of FILE or CENTERA can appear to be mounted more than once.
3. In the message ANR8448I, the drive name is listed as UNKNOWN for volumes of the FILE device type with a non-shared device class. The reason is that no drive is associated with the volumes; drive names are shown in the file-based library.
4. If you issue the `QUERY MOUNT` command while the drive is being cleaned, the command output continues to show a DISMOUNTING status for the dismounted volume until the cleaning completes.

Example: Display detailed information about mounted sequential volumes

Display details about mounted volumes.

```
query mount format=detailed

ANR2017I Administrator SERVER_CONSOLE issued command: QUERY
MOUNT format=detailed
ANR8487I Mount point in device class FILE is waiting for the
volume mount to
complete -- owning server: SERVER1, status: WAITING FOR VOLUME
(session: 0, process: 1).
ANR8488I LTO volume 015005L4 is mounted R/W in drive IBMVTL1
(/dev/rmt37) -- owning
server: SERVER1, status: IN USE (session: 0, process: 2).
ANR8486I Mount point in device class FILE is reserved -- owning
server: SERVER1,
status: RESERVED (session: 5, process: 0).
ANR8334I 3 matches found.
```

Related commands

Table 1. Commands related to QUERY MOUNT

Command	Description
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
REPLY	Allows a request to continue processing.

QUERY NASBACKUP (Query NAS backup images)

Use this command to display information about the file system image objects that have been backed up for a specific NAS node and file space. You can only use this command to display objects that were backed up for a NAS node using NDMP.

The server displays all matching objects, the dates that these objects were backed up, and information about a table of contents (TOC) for the object.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query NASBackup--node_name--fileSPACE_name----->
. -BEGINDate-----TODAY - 7-. . -BEGINTime-----00:00:00-.
>--+-----+-----+-----+-----+----->
' -BEGINDate-----date-----' ' -BEGINTime-----time-----'

. -ENDDate-----TODAY-. . -ENDTime-----23:59:59-.
>--+-----+-----+-----+-----+----->
' -ENDDate-----date--' ' -ENDTime-----time-----'

. -TYPE-----BACKUPImage----- .
>--+-----+-----+-----+-----+----->>
' -TYPE-----+BACKUPImage+-'
' -SNAPMirror--'
```

Parameters

node_name (Required)

Specifies the name of the NAS node for which backup objects are displayed. You cannot use wildcards to specify this name.

fileSPACE_name (Required)

Specifies the name of the file space for which backup objects are displayed. You can use wildcards to specify this name.

BEGINDate

Specifies the beginning date to select the backup objects to display. All backup objects that were created on or after the specified date are displayed. The default is seven days prior to the current date. You can use this parameter with the **BEGINTIME** parameter to specify a range for the date and time. This parameter is optional.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/2002
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To display information about the image objects that have been created a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE= -7 .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time to select the backup objects to display. All backup objects created on or after the specified time display. This parameter is optional. The default is midnight (00:00:00) on the date specified for the **BEGINDATE**.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 <i>or</i> +03:00. If you issue this command at 9:00 with <code>BEGINTIME=NOW+3</code> or <code>BEGINTIME=+3</code> , the server displays image objects with a time of 12:00 or later on the begin date.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-04:00 <i>or</i> -04:00. If you issue this command at 9:00 with <code>BEGINTime=NOW-3:30</code> or <code>BEGINTime= -3:30</code> , the server displays image objects with a time of 5:30 or later on the begin date.

ENDDate

Specifies the ending date used to select the backup objects to be displayed. All backup objects created on or before the specified date are displayed. This parameter is optional. The default is the current date. You can use this parameter with the `ENDTIME` parameter to specify an ending date and time.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/2002
TODAY	The current date	TODAY
TODAY-days <i>or</i> -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 <i>or</i> -1. To display information created up to yesterday, you can specify <code>ENDDATE=TODAY-1</code> or simply <code>ENDDATE= -1</code> .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time used to select the backup objects to be displayed. All backup objects created on or before the specified time are displayed. This parameter is optional. The default is 23:59:59. You can use this parameter with the `ENDDATE` parameter to specify a range for the date and time.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, the server displays image objects with a time of 12:00 or later on the end date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, the server displays image objects with a time of 5:30 or later on the end date you specify.

TYPE

Specifies the type of NDMP backup images for which you want to display information. The default value for this parameter is BACKUPIIMAGE. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

BACKUPIImage

Specifies that the output should show only the standard NAS base and differential images. This is the default value for this parameter.

SNAPMirror

Specifies whether to display information about NetApp SnapMirror images. SnapMirror images are block-level full-backup images of a file system. A SnapMirror image can only be restored to a file system that has been prepared as a SnapMirror target volume. Refer to the documentation that came with your NetApp file server for more information. This parameter is valid for NetApp and IBM N-Series file servers only.

Example:

Issue the QUERY NASBACKUP command to display information about a node, nas1, and a filesystem, /vol/vol1.

```
query nasbackup nas1 /vol/vol1
```

Node Name	Filespace Name	Object Type	Object Size (MB)	Creation Date	Has Contents	Mgmt Class	Image Storage Pool Name
NAS1	vol/vol1	Full image	1050.5	10/22/2002 10:50:57	YES	DEFAULT	NASBACKUPS
NAS1	vol/vol1	Differential image	9.1	10/22/2002 11:03:21	YES	DEFAULT	NASBACKUPS
NAS1	vol/vol1	Full image	1050.5	10/22/2006 10:43:00	YES	STANDARD	FILEPOOL
NAS1	vol/vol1	Differential image	9.1	10/25/2006 11:53:21	YES	STANDARD	FILEPOOL

Example:

Issue the QUERY NASBACKUP command to display information about all NetApp SnapMirror to Tape images for a node, nas2, and a filesystem, /vol/vol2.

```
query nasbackup nas2 /vol/vol2 type=snapmirror
```

Node Name	Filespace Name	Object Type	Object Size (MB)	Creation Date	Mgmt Class	Image Storage Pool Name
NAS2	vol/vol2	SnapMirror	1050.5	04/02/2008 10:50:57	STANDARD	MYPOOL
NAS2	vol/vol2	SnapMirror	1450.5	04/02/2008 11:03:21	STANDARD	MYPOOL

Field descriptions

Domain

Specifies a list of policy domains that limit the client node query. Only nodes that are assigned to one of the specified policy domains are displayed. This parameter is optional. Separate the items in the list by commas, with no intervening spaces. You can use wildcard characters to specify a domain. All clients that are assigned to a matching domain are displayed. If you do not specify a value for this parameter, all policy domains are included in the query.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. You can specify one of the following values:

Standard

Specifies that partial information is displayed for the specified client nodes.

Detailed

Specifies that complete information is displayed for the specified client nodes.

Type

Specifies the type of node to include in the query results. The parameter is optional. The default value is CLIENT. You can specify one of the following values:

Any

Specifies any type of node.

Client

Specifies client nodes that are backup-archive clients, IBM Spectrum Protect™ for Space Management clients, or application clients.

NAS

Specifies NAS nodes.

Server

Specifies client nodes that are other servers.

Authentication

Specifies the password authentication method for the node.

Local

Display those nodes that authenticate to the IBM Spectrum Protect server.

LDap

Display those nodes that authenticate to an LDAP directory server. The node password is case-sensitive.

Example: Display information about registered client nodes

Display information about all registered client nodes.

```
query node
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
CLIENT1	AIX	STANDARD	6	6	No
GEORGE	AIX	STANDARD	1	1	No
JANET	AIX	STANDARD	1	1	No
JARED	Linux86	STANDARD	1	1	No
JOE2	Mac	STANDARD	<1	<1	No
TOMC	WinNT	STANDARD	1	1	No

Example: Displayed detailed information about a client node

Display complete information about the client node named Joe.

```
query node joe format=detailed
```

```
Node Name: JOE
Platform: WinNT
Client OS Level: 4.00
Client Version: Version 5, Release 4,
Level 0.0
Application Version: Version 6, Release 4,
Level 0.4
```

```

Policy Domain Name: STANDARD
Last Access Date/Time: 09/24/2012 18:55:46
Days Since Last Access: 6
Password Set Date/Time: 09/24/2012 18:26:43
Days Since Password Set: 6
Invalid Sign-on Count: 0
Locked?: No
Contact:
Compression: Client
Archive Delete Allowed?: Yes
Backup Delete Allowed?: No
Registration Date/Time: 09/24/2012 18:26:43
Registering Administrator: SERVER_CONSOLE
Last Communication Method Used: Tcp/Ip
Bytes Received Last Session: 108,731
Bytes Sent Last Session: 698
Duration of Last Session: 0.00
Pct. Idle Wait Last Session: 0.00
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Optionset:
URL: http://joe.host.name:1581
Node Type: Client
Password Expiration Period: 60
Keep Mount Point?: No
Maximum Mount Points Allowed: 2
Auto Filespace Rename: No
Validate Protocol: No
TCP/IP Name:
TCP/IP Address: 9.11.153.39
Globally Unique ID: 11.9c.54.e0.8a.b5.11.d6.b3.
c3.00.06.29.45.c1
Transaction Group Max: 0
Data Write Path: ANY
Data Read Path: ANY
Session Initiation: ClientOrServer
High-level Address:
Low-level Address: 1501
Collocation Group Name:
Proxynode Target:
Proxynode Agent:
Node Groups:
Email Address:
Deduplication: ServerOnly

```

AIX | Linux

```

Users allowed to back up: ALL
Replication State: Enabled
Replication Mode: Send
Backup Replication Rule: DEFAULT
Archive Replication Rule: ALL_DATA
Space Management Replication Rule: None
Replication Primary Server: PRODSERVER1
Last Replicated to Server: DRSERVER1
Client OS Name: WIN: Windows XP
Client Processor Architecture: x86
Client Products Installed: WIN, FCM, VE
Client Target Version:
Version 6, Release 2, Level 0.0
Authentication: Local
SSL Required: No
Session Security: Strict
Transport Method: TLS 1.2
Split Large Objects: Yes
At-risk type: Default interval
At-risk interval:
Utility URL:
Replication Recovery of Damaged Files: Yes
Decommissioned:
Decommissioned Date:

```

Field descriptions

Node Name

The name of the client node.

Platform

The operating system of the client node, as of the last time that the client node contacted the server. A question mark (?) is displayed until the client node first accesses the server and reports its operating system type.

Client OS Level

The level of the operating system for the client as of the last time that the client node contacted the server.

Client Version

The version of the client that is installed on the client node.

This field does not apply to NAS nodes.

Application Version

The version of the Data Protection for VMware client.

Policy Domain Name

The assigned policy domain of the client node.

Last Access Date/Time

The last date and time that the client node accessed the server.

Days Since Last Access

The number of days that elapsed since the last time that the client node accessed the server.

Password Set Date/Time

The date and time that the password was set for the client node.

Days Since Password Set

The number of days that elapsed since the password was set for the client node.

Invalid Sign-on Count

The number of invalid sign-on attempts that were made since the last successful sign-on. This count can be non-zero only when the invalid password limit (SET INVALIDPWLIMIT) is greater than zero. When the number of invalid attempts equals the limit that is set by the SET INVALIDPWLIMIT command, the node is locked out of the system.

Locked?

Whether the client node is locked out of IBM Spectrum Protect.

Contact

Any contact information for the client node.

Compression

Whether compression is enabled on the client node.

This field does not apply to NAS nodes.

Archive Delete Allowed?

Whether the client node can delete its own archive files.

Backup Delete Allowed?

Whether the client node can delete its own backup files.

Registration Date/Time

The date and time that the client node was registered.

Registering Administrator

The name of the administrator that registered the client node.

Last Communication Method Used

The communication method that was last used by the client node to contact the server.

Bytes Received Last Session

The number of bytes received by the server during the last client node session.

This field does not apply to NAS nodes.

Bytes Sent Last Session

The number of bytes sent to the client node.

This field does not apply to NAS nodes.

Duration of Last Session

How long the most recent client node session lasted, in seconds.

This field does not apply to NAS nodes.

Pct. Idle Wait Last Session

The percentage of the total session time that the client was not running any functions.

This field does not apply to NAS nodes.

Pct. Comm. Wait Last Session

The percentage of the total session time that the client waited for a communication response from the server.

This field does not apply to NAS nodes.

Pct. Media Wait Last Session

The percentage of the total session time that the client waited for a removable volume to be mounted.

This field does not apply to NAS nodes.

Optionset

The name of the client option set.

URL

The URL of the IBM Spectrum Protect web client that is configured on the client system. You can use the URL in a web browser and in the Operations Center to remotely manage the client node.

Node Type

The type of client node. One of the following values is possible:

- Client: a backup-archive client, an IBM Spectrum Protect for Space Management client, or an application client
- Server: an IBM Spectrum Protect server
- NAS: a NAS file server

Password Expiration Period

The password expiration period of the client node.

Keep Mount Point?

Whether the client node retains a mount point during a session.

Maximum Mount Points Allowed

The number of mount points that a client node can use on the server for IBM Spectrum Protect for Space Management migration and for backup and archive operations. This parameter does not apply to nodes with a type of NAS or SERVER. If a client node was registered to a server at Version 3.7 or later, the value is 0-999, depending on the value that is set with the MAXNUMMP parameter of the REGISTER NODE command. If the client node was registered under previous versions of the server and the MAXNUMMP parameter was not explicitly set by using the UPDATE NODE command, the value is set to NOLIMIT. The MAXNUMMP value is not evaluated or enforced during client data read operations such as restore, retrieve, and IBM Spectrum Protect for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node. This evaluation might prevent the data store operations from acquiring mount points.

Auto Filespace Rename

Whether IBM Spectrum Protect prompts the client to rename file spaces when the client system upgrades to a client that supports Unicode. This field is valid only for client systems that use Windows, Macintosh OS X, or NetWare operating systems.

Validate Protocol (deprecated)

Whether the client has data validation enabled. If the client has data validation enabled, this field specifies whether IBM Spectrum Protect validates only the file data or all data, which includes file metadata. You can enable data validation by using the REGISTER NODE or UPDATE NODE command. This field is deprecated.

TCP/IP Name

The host name of the client node as of the last time that the client node contacted the server. The field is blank if the client software does not support reporting this information to the server.

TCP/IP Address

The TCP/IP address of the client node as of the last time that the client node contacted the server. The field is blank if the client software does not support reporting this information to the server.

Globally Unique ID

The globally unique identifier (GUID) as of the last time that the client node contacted the server. This GUID identifies the host computer on which the node is located.

Transaction Group Max

Specifies the number of files per transaction committed that are transferred between a client and a server. Client performance might be improved by using a larger value for this option.

Data Write Path

Specifies the transfer path that is used when the client sends data to the server, storage agent, or both, during storage operations. If a path is unavailable, the node cannot send any data.

AIX | **Linux** Data transfer path options are ANY, LAN, or LAN-free.

Data Read Path

Specifies the transfer path that is used when the server, storage agent, or both, read data for a client, during operations such as restore or retrieve. If a path is unavailable, data cannot be read.

AIX | **Linux** Data transfer path options are ANY, LAN, or LAN-free.

Session Initiation

Controls whether the server or client initiates sessions. The following two options are available:

- ClientOrServer
- Serveronly

High-level Address

Specifies the client IP address that the server contacts to initiate scheduled events when SESSIONINITIATION is set to SERVERONLY.

Low-level Address

Specifies the client port number on which the client listens for sessions from the server when SESSIONINITIATION is set to SERVERONLY.

Collocation Group Name

Specifies the name of the collocation group to which a node belongs. If a node does not belong to a collocation group, this field is blank.

Tip: If the node contains file spaces that are members of a file space collocation group, this field is left blank. You can find file space names by issuing the QUERY FILESPACE command.

Proxynode Target

Specifies which nodes are proxy nodes (agents) for other nodes, in a space-separated list. If there are no nodes in that type of association, this field is blank.

Proxynode Agent

Specifies the originating (target) node name for a proxy node session, in a space separated list. If there are no nodes in that type of association, this field is blank.

Node Groups

Specifies the name of the node group to which a node belongs. If a node does not belong to a node group, this field is blank.

Email Address

Specifies the email address of the client node.

Deduplication

The location where data is deduplicated. The value ServerOnly specifies that data stored by this node can be deduplicated on the server only. The Clientorserver value specifies that data stored by this node can be deduplicated on either the client or the server.

AIX | **Linux** Users allowed to back up

AIX | **Linux** Specifies whether a non-root user ID or only a root user ID can back up files to the server. ALL indicates all users, while ROOT indicates that just the root user ID can back up files to the server. This output is not available if the client node operating system is considered a single-user operating system.

Replication State

Indicates whether the node is enabled for replication. The following values are possible:

Enabled

The node is configured for replication and ready to replicate.

Disabled

The node is configured for replication but is not ready to replicate.

None

The node is not configured for replication.

Replication Mode

Indicates whether the node is configured as the source of or target for replicated data. If this field is blank, the node is not configured for replication. The following values are possible:

Send

The node is configured as the source of data for replication.

Receive

The node is configured as the target of data for replication.

SyncSend

The data that belongs to the node is to be synchronized with the node data that is on the target replication server. Synchronization applies only to nodes whose data was imported from a source replication server and imported to the target replication server. Synchronization occurs during replication.

SyncReceive

The data that belongs to the node is to be synchronized with the node data that is on the source replication server. Synchronization applies only to nodes whose data was imported from a source replication server and imported to the target replication server. Synchronization occurs during replication.

None

The node is not configured for replication.

Replication Primary Server

Specifies the source replication server for the client node.

Backup Replication Rule

Archive Replication Rule

Space Management Replication Rule

The replication rule that applies to back up, archive, and space-managed data that belongs to the node. The following values are possible:

ALL_DATA

Replicates backup, archive, or space-managed data. The data is replicated with normal priority.

ACTIVE_DATA

Replicates active backup data. The data is replicated with normal priority.

Attention: If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the `FORCERECONCILE=YES` parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Replicates backup, archive, or space-managed data. The data is replicated with high priority.

ACTIVE_DATA_HIGH_PRIORITY

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority.

DEFAULT

Replicates backup, archive, or space-managed data according to the domain rule for the data type.

NONE

No data is replicated. For example, if the replication rule for archive data is NONE, archive data that belongs to the node is not replicated.

Last Replicated to Server

Specifies the name of the server that the node was last replicated to and the name of the server that the client fails over to during restore operations.

Client OS Name

The operating system of the client. The client deployment wizard uses this information to deploy a package to the client.

This field is reported only for IBM Spectrum Protect clients at V6.2.0.0 and later.

Client Processor Architecture

The client architecture. The client deployment wizard uses this value to determine which package to deploy when the client is being updated. This field is reported only for IBM Spectrum Protect clients at V6.2.0.0 and later.

Client Products Installed

The products that are on the node. The following products might be listed:

- BA (Backup-Archive Client)
- VE (Virtual Environments)
- FCM (FlashCopy® Manager)

Client Target Version

The version of the client that is installed at a time that is scheduled through the DEFINE SCHEDULE or UPDATE SCHEDULE command. This field is reported only for IBM Spectrum Protect clients at V6.2.0.0 and later.

Authentication

Specifies the password authentication method: LOCAL, LDAP, or LDAP (pending).

Authentication Target	Authentication Method
-----------------------	-----------------------

Authentication Target	Authentication Method
IBM Spectrum Protect server	LOCAL
LDAP directory server	LDAP
This node is configured to authenticate with an LDAP directory server, but the node did not yet authenticate.	LDAP (pending)

SSL Required (deprecated)

Specifies whether the security setting for the node requires the Secure Sockets Layer (SSL) protocol. Values can be YES, NO, or Default. You must have system level authority to update the node SSLREQUIRED setting. This field is deprecated.

Session Security

Specifies the level of session security that is enforced for the node. Values can be STRICT or TRANSITIONAL.

Transport Method

Specifies the transport method that was last used for the specified node. Values can be TLS 1.2, TLS 1.1, or NONE. A question mark (?) is displayed until a successful authentication is completed.

Split Large Objects

Specifies whether large objects that are stored by this node are automatically split into smaller pieces, by the server, to optimize server processing. Yes indicates that the server splits large objects (over 10 GB) into smaller pieces when stored by a client node. No indicates that this process is bypassed. The default value is Yes.

At-risk type

Specifies the at-risk evaluation type. Values can be Default, Bypassed, or Custom. Default indicates that the node is evaluated with the same interval that was specified for the nodes classification by the SET STATUSATRISKINTERVAL command. Bypassed indicates that the node is not evaluated for at-risk status by the status monitor. Custom indicates that the node is evaluated with the interval that was specified by the SET NODEATRISKINTERVAL command, rather than the interval that was specified by the SET STATUSATRISKINTERVAL command.

At-risk interval

Specifies the number of hours between two client backup activities, or two replication activities, after which the status monitor indicates that the activity is at risk. This field contains a value only when the *At-risk type* field contains the value of *Custom*.

Utility URL

Specifies the address of the IBM Spectrum Protect client management services that are configured on the client system. This URL is used by the Operations Center to access client log files so that you can remotely diagnose client issues from the Operations Center.

Replication Recovery of Damaged Files

Specifies whether damaged files can be recovered for this node from a target replication server.

Decommissioned

Specifies whether the client node is decommissioned. The following values are possible:

YES

Specifies that the node is decommissioned.

Null value

Specifies that the node is not decommissioned.

PENDING

Specifies that the node is being decommissioned, or the decommission process failed.

Tip: If you want to determine the status of a pending decommission process, follow the instructions in Decommissioning a client node.

Decommissioned Date

Specifies the date that the client node was decommissioned.

Example: Display information about node roles

The example output is only a portion of the full display.

```
query node alvin f=d
```

```

    Proxynode Agent:
      Node Groups:
      Email Address:
      Deduplication: ServerOnly
    Users allowed to back up: All
```

```

Role: Server
Role Override: UseReported
Processor Vendor: ORACLE
Processor Brand: UltraSPARC-T2
Processor Type: 4
Processor Model:
Processor Count: 1
Hypervisor:
API Application: NO
Scan Error: NO
MAC Address:

```

Field Descriptions

Role

The processor role as reported by the client.

Role Override

The override value for role, which is specified with the UPDATE NODE command.

Processor Vendor

The processor vendor as reported by the client.

Processor Brand

The processor brand as reported by the client.

Processor Type

The processor type as reported by the client. This value specifies the number of processor cores that are used for PVU calculation.

Processor Model

The processor model as reported by the client.

Processor Count

The processor count as reported by the client.

Hypervisor

The hypervisor as reported by the client.

API Application

The client indicator that the client is an API application.

Scan Error

The indicator of whether the latest scan for processor information might be failing and needs investigation.

MAC Address

MAC Address as reported by the client.

Example: View all nodes that authenticate to the IBM Spectrum Protect server

If you want to view all nodes that authenticate locally, specify the following command:

```
query node * authentication=local
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
NODE1	WinNT	STANDARD	3	3	No
LOCAL	(?)	STANDARD	7	7	No

Related commands

Table 1. Commands related to QUERY NODE

Command	Description
LOCK NODE	Prevents a client from accessing the server.
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

Command	Description
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REMOVE REPLNODE	Removes a node from replication.
RENAME NODE	Changes the name for a client node.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.
SET MINPWLENGTH	Sets the minimum length for client passwords.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.
UPDATE NODE	Changes the attributes that are associated with a client node.

QUERY NODEDATA (Query client data in volumes)

Use this command to display information about the data for one or more nodes in a sequential access storage pool. QUERY NODEDATA displays the name of the volume on which a node's data is written and the amount of space that is occupied by the data on that volume. This information is useful when you determine how to group nodes into collocated storage pools.

Privilege class

Restriction: You cannot use this command to display information for container storage pools.

Any administrator can issue this command.

Syntax

```

      .-.-.-.-.-.
      v          |
>>-Query NODEData--+-node_name-+----->
                    '-COLLOCGroup--colloc_group-'
>--+-----+-----+-----><
    '-STGpool---pool_name-'  '-VOLUME---vol_name-'

```

Parameters

node_name

Specifies the name of the client node for which you want to locate data. You can specify one or more names. If you specify multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple names. You must specify either a node name or collocation group name, but not both.

COLLOCGroup

Specifies the name of the collocation group for which you want to locate data. You must specify either a node name or collocation group name, but not both.

Important: If the amount of space that is needed to complete the query about a collocation group exceeds the SQL buffer limit, the QUERY NODEDATA command can fail. If the command fails for this reason, issue the QUERY COLLOCGROUP command to display a list of nodes in the group. Then, issue the QUERY NODEDATA command for each node in the group.

STGpool

Specifies the name of the sequential storage pool to query. This parameter is optional. You can use wildcard characters to specify the names. If a wildcard matches the name of a disk storage pool, the name of the disk storage pool is ignored. If you do not specify a value for this parameter, all sequential storage pools are queried.

VOLume

Specifies the volume that contains the data. This parameter is optional. You can use wildcard characters to specify multiple names. If you do not specify a value for this parameter, all volumes in the storage pool are queried.

Use wildcards to display node data for a sequential access storage pool

Display information about where node data is stored in a sequential storage pool. Use a wildcard character to indicate node names. See Field descriptions for field descriptions.

```
query nodedata e*
```

Node Name	Volume Name	Storage Pool Name	Physical Space Occupied (MB)
EDU_J2	E:\tsm\server\00000117.BFS	EDU512	0.01
EDU_J2	E:\tsm\server\00000122.BFS	EDU319	0.01
EDU_J3	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_J3	E:\tsm\server\00000120.BFS	EDU319	0.01
EDU_J7	E:\tsm\server\00000118.BFS	EDU512	0.04
EDU_J7	E:\tsm\server\00000123.BFS	EDU319	0.04
EDU_JJ1	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_JJ1	E:\tsm\server\00000121.BFS	EDU512	0.01

Display node data information for a specific collocation group

Display information about the location of node data in a sequential storage pool for a particular collocation group. In this example, nodes EDU_J3 and EDU_JJ1 are the only members that belong to collocation group, grp1, and have data in a sequential access storage pool.

```
query nodedata collocgroup=grp1
```

Node Name	Volume Name	Storage Pool Name	Physical Space Occupied (MB)
EDU_J3	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_J3	E:\tsm\server\00000120.BFS	EDU319	0.01
EDU_JJ1	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_JJ1	E:\tsm\server\00000121.BFS	EDU512	0.01

If you specify a file space collocation group, only the volumes of the file spaces that belong to the collocation group are displayed. If you specify a file space collocation group and a volume, the file space volumes within the collocation group that are also in the specified volume are displayed.

Field descriptions

Node Name

Specifies the name of the node.

Volume Name

Specifies the name of the volume that contains the node data.

Storage Pool Name

Specifies the name of the storage pool in which the volume is located.

Physical Space Occupied (MB)

Specifies the amount of physical space that is occupied by the node's data. Physical space includes empty space within aggregates, from which files might be deleted or expired.

Related commands

Table 1. Commands related to QUERY NODEDATA

Command	Description
---------	-------------

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

QUERY NODEGROUP (Query a node group)

Use this command to display the node groups defined on the server.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*-----
>>-Query NODEGroup--+-----+----->
      '-group_name-'

      .-Format---Standard----.
>--+-----+-----<<
      '-Format---+Standard-+-'
      '-Detailed-'

```

Parameters

group_name

Specifies the name of the node group to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all node groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed. To display the members of the node group, you must specify FORMAT=DETAILED.

Example: List node groups on the server

Display the node groups defined on the server. See Field descriptions for field descriptions.

```
query nodegroup
```

Node Group Name	Node Group Description
DEPT_ED	Education department
GROUP1	Low cap client nodes.

Example: Display detailed node group information

Display complete information about all node groups and determine which client nodes belong to which node groups. See Field descriptions for field descriptions.

```
query nodegroup format=detailed
```

```
Node Group Name: DEPT_ED
Node Group Description: Education department
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 04/21/2006 10:59:03
Node Group Member(s): EDU_1 EDU_7

Node Group Name: GROUP1
Node Group Description: Low cap client nodes.
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 04/21/2006 10:59:16
Node Group Member(s): CHESTER REX NOAH JARED
```

Field descriptions

Node Group Name

The name of the node group.

Node Group Description

The description for the node group.

Last Update by (administrator)

The name of the administrator that defined or most recently updated the node group.

Last Update Date/Time

The date and time that an administrator defined or most recently updated the node group.

Node Group Member(s)

The members of the node group.

Related commands

Table 1. Commands related to QUERY NODEGROUP

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

QUERY OCCUPANCY (Query client file spaces in storage pools)

Use this command to show where client file spaces are stored and how much space they occupy.

Specifies that archive files are queried.

SPacem

Specifies that space-managed files (files that were migrated by an IBM Spectrum Protect for Space Management client) are queried.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with Unicode support. A backup-archive client with Unicode support is available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare. Use this parameter only when you specify a partly or fully qualified file space name.

The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the names and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specifies how you want the server to interpret the file space names that you enter. Use this parameter only when you enter a single wildcard character for the file space name or when you do not specify any file space name.

The default value is BOTH, which means that the file spaces are included regardless of code page type. Possible values are:

UNICODE

Include file spaces that are only Unicode enabled.

NONUNICODE

Include file spaces that are not only Unicode enabled.

BOTH

Include file spaces regardless of code page type.

Example: Display file spaces assigned to a specific node

Display information about where all file spaces assigned to the node named DAISY are stored. See Field descriptions for field descriptions.

```
query occupancy daisy
```

Node Name	Type	Filespace Name	FSID	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
DAISY	Bkup	DRIVED	1	COPYFILE	38	0.45	0.42

Example: Display file spaces assigned to a specific node with a backup file type

Display information about the file spaces that belong to the node WAYNE, and that have a backup file type. See Field descriptions for field descriptions.

```
query occupancy wayne type=backup
```

Node Name	Type	Filespace Name	FSID	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
WAYNE	Bkup	DWG1	1	BACKUPPOOL1	2,330	53.19	50.01
WAYNE	Bkup	OS2C	2	BACKUPPOOL1	1,554	32.00	31.30

Field descriptions

Node Name

The node that owns the file space. If the node was previously deleted, the node name DELETED is displayed.

Type

The type of data. Possible values are:

Arch

Data that has been archived.

Bkup

Data that has been backed up.

SpMg

Data that has been migrated from an IBM Spectrum Protect for Space Management client.

Filespace Name

The name of the file space that belongs to the node.

If the file space was previously deleted, the file space name DELETED is displayed.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Storage Pool Name

The storage pool where the file space is located.

Number of Files

The number of logical files that belong to the file space and are stored in this storage pool. When storing a file larger than 10 GB, the server splits the file into 10 GB fragments. The number of fragments is also included in this value for occupancy calculations.

Physical Space Occupied (MB)

The amount of physical space that is occupied by the file space. Physical space includes empty space within aggregates, from which files might have been deleted or expired. For this value, 1 MB = 1048576 bytes.

Tip: This field does not display a value for storage pools that are set up for data deduplication. If you turn off data deduplication for a storage pool, a value for physical occupancy is not displayed until the storage pool is empty of deduplicated files.

Logical Space Occupied (MB)

The amount of space that is occupied by logical files in the file space. Logical space is the space that is actually used to store files, excluding empty space within aggregates. For this value, 1 MB = 1048576 bytes.

FSID

The file space ID (FSID) for the file space. The server assigns a unique FSID when a file space is first stored on the server.

Related commands

Table 1. Commands related to QUERY OCCUPANCY

Command	Description
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.

QUERY OPTION (Query server options)

Use this command to display information about server options.

Change server options by editing the server options file or by issuing the SETOPT command. When you edit the server options file, you must restart the server before any changes take effect. Any changes you make by issuing the SETOPT command take effect immediately.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query OPTion--+-*-----+----->>
                    '-optionname-'
```

Parameters

optionname

Specifies the name of an option in the server options file. This parameter is optional. You can use wildcard characters to specify this name. All matching server options display. If you do not specify this parameter, information on all options displays.

Example: Display all server options

Display general information about all server options. The output lists all options with their specified values.

```
query option
```

Example: Display options settings using a wildcard character

View the option settings for all options that begin with L.

```
query option l*
```

Server Option	Option Setting
Language	AMENG

Example: Display LDAP directory servers

View the settings for all LDAP directory servers.

```
query option ldapurl
```

Server Option	Option Setting
LDAP URL	ldap:\\tophoy.tucson.com\cn=tsmdata
LDAP URL	ldap:\\krypton.ibm.com\ou=tsmdata,dc=ibm,dc=com

Field descriptions

Server Option

Specifies the name of the option in the server options file.

Option Setting

Specifies the name of the option in the server options file.

Related commands

Table 1. Commands related to QUERY OPTION

Command	Description
SETOPT	Updates a server option without stopping and restarting the server.

QUERY PATH (Display a path definition)

Use this command to display the path between a source and a destination.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query PATH-----+----->
      | .-*-----+-----|
      |'-source_name-----+'
      |'-destination_name-'
      |
      |.-SRCType----ANY-----+----->
      |'-SRCType----ANY-----+'
      |'+DATAMover+'
      |'-SERVer----+'
      |
      |.-DESTType----ANY-----+----->
      |'-DESTType----ANY-----+'
      |'+DRIVE--LIBRARY----library_name+'
      |'-LIBRARY-----+'
      |
      |.-Format----Standard----+-----><
      |'-Format----Standard----+'
      |'-Detailed-'
```

Parameters

source_name

Specifies the name of a source for which to display paths. This parameter is optional. You can specify wildcard characters. The default is to display paths for all sources.

A source is a data mover, a server, or a storage agent.

destination_name

Specifies the name of a destination for which to display paths. This parameter is optional. You can specify wildcard characters. The default is to display paths for all destinations.

SRCType

Specifies the type of the source. This parameter is optional. The default is to display paths for all source types. Possible values are:

ANY

Specifies to display paths with any source type.

DATAMover

Specifies to only display paths with the DATAMOVER source type.

SERVer

Specifies to only display paths with the SERVER source type. (A source that has a source type of SERVER is a storage agent.)

DESTType

Specifies the type of the destination. This parameter is optional. The default is to display paths for all destination types. Possible values are:

ANY

Specifies to display paths with any destination type.

DRive

Specifies to display only paths with the DRIVE destination type. When the destination type is a drive, you must specify the library name. You can refine which paths are displayed by entering a name in the LIBRARY parameter.

LIBRARY

Specifies that only paths with destination type LIBRARY display.

LIBRARY

Specifies the name of the library to which the drive belongs. This parameter is required when the destination type is a drive (DESTTYPE=DRIVE).

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary path information

Display information about paths for the source NETAPP1. See Field descriptions for field descriptions.

```
query path netapp1
```

Source Name	Source Type	Destination Name	Destination Type	Online
NETAPP1	DATAMOVER	DRIVE1	DRIVE	Yes
NETAPP1	DATAMOVER	NASLIB	LIBRARY	Yes

Example: Display detailed path information

Display detailed information about paths for the source NETAPP1. See Field descriptions for field descriptions.

```
query path netapp1 format=detailed
```

Linux

```
Source Name: NETAPP1
Source Type: DATAMOVER
Destination Name: NASLIB
Destination Type: LIBRARY
Library:
Device: /dev/tmsmcsi/mc0
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2002 20:52:56
```

```
Source Name: NETAPP1
Source Type: DATAMOVER
Destination Name: DRIVE1
Destination Type: DRIVE
Library: NASLIB
Device: rst01
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2002 20:55:23
```

AIX

Windows

```
Source Name: NETAPP1
Source Type: DATAMOVER
Destination Name: NASLIB
Destination Type: LIBRARY
Library:
Device: mc0
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2001 20:52:56
```

```
Source Name: NETAPP1
```

```

Source Type: DATAMOVER
Destination Name: DRIVE1
Destination Type: DRIVE
Library: NASLIB
Device: rst01
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2001 20:55:23

```

AIX Linux

Example: Display detailed path information for a z/OS media server

Display detailed information about a z/OS® media server path. See Field descriptions for field descriptions.

```

query path format=detailed

Source Name: SERVER1
Source Type: SERVER
Destination Name: ZOSMEDIA
Destination Type: LIBRARY
Library:
Node Name:
Device:
External Manager:
ZOS Media Server: MEDSERV1
Comm. Method:
LUN:
Initiator: 0
Directory:
On-Line: Yes
Last Update by (administrator): ADMIN
Last Update Date/Time: 06/08/2011 15:33:39

```

Field descriptions

Source Name

The name of the source.

Destination Name

The name of the destination.

Source Type

The type of the source.

Destination Type

The type of the destination.

Library

The name of the library that contains the drive that is the destination.

This field will be blank if the destination type is library. The library name is in destination name field when the destination is a library.

Node Name

The name of the device that is the destination.

Device

The name of the device that is the destination.

External Manager

The name of the external manager.

ZOS Media Server

The name of the z/OS media server.

Comm. Method

Specifies the type of communication method.

LUN

Specifies the logical unit name through which the disk can be accessed by the source.

Initiator

Specifies the initiator of the communication.

Directory

Specifies the directory location of a file on the source.

Policy Domain Name	Policy Set Name	Default Mgmt Class Name	Description
EMPLOYEE-RECORDS	ACTIVE	ACTIVEFILES	Personnel Department
EMPLOYEE-RECORDS	HOLIDAY	ACTIVEFILES	Personnel Department
EMPLOYEE-RECORDS	VACATION	ACTIVEFILES	Personnel Department
PROG1	SUMMER		Programming Group Policies
PROG2	SUMMER		Programming Group Policies
STANDARD	ACTIVE	STANDARD	Installed default policy set.
STANDARD	STANDARD	STANDARD	Installed default policy set.

Example: Displayed detailed information about a specific policy set

Query the VACATION policy set that is in the EMPLOYEE_RECORDS policy domain. Create the output in detailed format. See Field descriptions for field descriptions.

```
query policyset employee_records vacation
format=detailed

Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Default Mgmt Class Name: ACTIVEFILES
Description: Personnel Department
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 05/31/1998 13:15:50
Managing profile: ADSM_INFO
Changes Pending: Yes
```

Field descriptions

Policy Domain Name

The name of the policy domain.

Policy Set Name

The name of the policy set.

Default Mgmt Class Name

The management class assigned as the default for the policy set.

Description

The description of the policy set.

Last Update by (administrator)

The name of the administrator or server that most recently updated the policy set. If this field contains \$\$CONFIG_MANAGER\$\$, the policy set is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

The date and time when the policy set was most recently defined or updated.

Managing Profile

The profile or profiles that manage the domain to which this policy set belongs.

Changes Pending

Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 1. Commands related to QUERY POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.

Command	Description
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY DOMAIN	Displays information about policy domains.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

QUERY PROCESS (Query one or more server processes)

Use this command to display information about active background processes.

To cancel background processes, issue the CANCEL PROCESS command. To display detailed information about node replication processes, issue the QUERY REPLICATION command.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query PProcess--+-+-----+----->
                        '-process_number-'
>+-----+-----+-----><
  '-DEscription----string-'  '-STATus----string-'
```

Parameters

process_number

Specifies the number of the background process to be queried. This parameter is optional. If not specified, information about all background processes is displayed.

DEscription

Specifies a text string that you want to search for in the list of active processes' descriptions. Enclose the string expression in quotation marks if it contains blanks. You can use text and a wildcard character to specify this string. This parameter is optional.

STATus

Specifies a text string that you want to search for in the list of active processes' statuses. Enclose the string expression in quotation marks if it contains blanks. You can use text and a wildcard character to specify this string. This parameter is optional.

Example: Query a single background process

Display information about background process 202. See Field descriptions for field descriptions.

```
query process 202
```

```
Process      Process      Process
Number      Description  Status
-----
      202  EXPORT SERVER  ANRONNNI EXPORT
Identifier MYEXPORTSERVER
ANR0648I Have copied the
following: 8 Domains 2
Policy Sets 10 Management
Classes 4 Copy Groups 1
Administrators 746 Bytes
(0 errors have been
detected) Current input
volume(s): C:\BUILD\540\
```


Example: Query all background processes

Display information about all background processes. See Field descriptions for field descriptions.

```
query process
```

Process Number	Process Description	Process Status
304	IDENTIFY DUPLICATES	Storage Pool FILEPOOL, Volume /tsmpool2/00006664. BFS, Files Processed: 2000, Duplicate Extents Found: 344, Duplicate Bytes Found: 3,238,123, Current Physical File (bytes): 2,626,676,296. Status: Processing
284	IDENTIFY DUPLICATES	Storage Pool FILEPOOL, Volume /tsmpool2/00006666. BFS, Files Processed: 2000, Duplicate Extents Found: 344, Duplicate Bytes Found: 3,238,123, Current Physical File (bytes): None. Status: Idle
4	Replicate Node	Replicating Node(s) IRONMAN. File spaces complete: 0. File spaces identifying and replicating: 1. File spaces replicating: 0. File spaces not started: 3. Files current: 11,920. Files replicated: 0 of 0. Files updated: 0 of 0. Files deleted: 0 of 0. Amount Replicated: 11,482 KB of 11,482 KB. Amount transferred: 11,482 KB. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).
37	Expiration	Processed 12 nodes out of 30 total nodes, examined 411 objects, deleting 411 backup objects, 0 archive objects, 0 DB backup volumes, 0 recovery plan files; 0 objects have been retried and 0 errors encountered.

Example: Query all background replication processes

Display information about all background replication processes. See Field descriptions for field descriptions.

```
query process desc="replicate node"
```

Process Number	Process Description	Process Status
4	Replicate Node	Replicating Node(s) IRONMAN. File spaces complete: 0. File spaces identifying and replicating: 1. File spaces replicating: 0. File spaces not started: 3. Files current: 11,920. Files replicated: 0

```

of 0. Files updated: 0 of 0.
Files deleted: 0 of 0. Amount
Replicated: 11,482 KB of 11,482
KB. Amount transferred: 11,482 KB.
Elapsed time: 0 Day(s), 0 Hour(s),
1 Minute(s).

```

Example: Query all background replication processes for a specific node

Display information about all background replication processes. See Field descriptions for field descriptions.

```
query process desc="replicate node" status=ironman
```

Process Number	Process Description	Process Status
4	Replicate Node	Replicating Node(s) IRONMAN. File spaces complete: 0. File spaces identifying and replicating: 1. File spaces replicating: 0. File spaces not started: 3. Files current: 11,920. Files replicated: 0 of 0. Files updated: 0 of 0. Files deleted: 0 of 0. Amount Replicated: 11,482 KB of 11,482 KB. Amount transferred: 11,482 KB. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).

Example: Verify that a replication recovery process was initiated

After you start a node replication process with file recovery enabled, verify that the target replication server initiated the file recovery process. Issue the QUERY PROCESS command on the target replication server. For descriptions of fields, see Field descriptions.

```
query process
```

Process Number	Process Description	Process Status
4	Replicate Node - Recovery.	Replicating node(s) 3MAUTOIMPORT. File spaces complete: 87. File spaces identifying and replicating: 0. File spaces replicating: 6. File spaces not started: 0. Files current: 0. Files replicated: 0 of 14. Files updated: 0 of 0. Files deleted: 0 of 0. Amount replicated: 0 KB of 11,688 bytes. Amount transferred: 0 KB. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).

Example: Verify that damaged files are being recovered during a replication process

After you start a node replication process with file recovery enabled, verify that damaged files are being recovered. Issue the QUERY PROCESS command on the source replication server. For descriptions of fields, see Field descriptions.

```
query process
```

Process Number	Process Description	Process Status
6	Replicate Node (As Secondary Recovery)	Recovering damaged files from server SERVER2, process 4, number of active sessions 10.

Example: Verify that the files are being converted

After you start a storage pool conversion process, verify that the files are being converted. For descriptions of fields, see Field descriptions.

```
query process
```

Process Number	Process Description	Process Status
6	Convert Stgpool	Converting storage pool FILEPOOL1 to directory-container storage pool NEWDEDUP1. Volumes Converted: 1 of 6, Volumes Failed: 0, Converted Files: 975, Converted Bytes: 196.27 MB, Skipped Files: 0, Skipped Bytes: 0 B, Total Bytes Transferred: 151.27 MB
7	Convert Stgpool	Converting storage pool DEDUPPOOL to directory-container storage pool DIRPOOL. Converted Files: 150 of 360, Converted Bytes: 79,598 KB of 388 MB. Unconverted Files: 12. Unconverted Bytes: 27 MB. Current input volume: /fvt/srv/BK01. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).
8	Convert Stgpool	Converting storage pool FILEPOOL1 to directory-container storage pool NEWDEDUP1. Converted Files: 0, Converted Bytes: 0 B of 1.00 GB, Skipped Files: 0, Skipped Bytes: 0 B, Total Bytes Transferred: 0 B, Current input volume: /STORAGE/file1/00000005.BFS, Elapsed time: 0 Days, 0 Hours, 1 Minutes.
10	Convert Stgpool	Converting storage pool FILEPOOL1 to directory-container storage pool NEWDEDUP1. Converted Files: 1007, Converted Bytes: 285.44 MB of 1.33 GB, Skipped Files: 0, Skipped Bytes: 0 B, Total Bytes Transferred: 196.28 MB, Current input volume: /STORAGE/file1/00000004.BFS, Elapsed time: 0 Days, 0 Hours, 1 Minutes.

Example: Verify movement from local disk to the cloud

After the data-transfer operation from the local disk to the cloud starts, verify that the data is moving. For descriptions of fields, see Field descriptions.

```
query process
```

Process Number	Process Description	Process Status
4	Local to Cloud Transfer	Local disk to cloud transfer for directory-container storage pool CLOUDPOOL. 1 container(s) processed. 2,100 KB in 4 data extent(s) transferred. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).

Field descriptions

Process Number

Specifies the number that is assigned to the active background process.

Process Description

Specifies a description of the active background process.

Process Status

Specifies the status of the active background process.

Tip: When a node replication process is finished on the target replication server, only end process information is stored in the activity summary table. The full summary for the replication process is stored in the activity summary table on the source replication server.

Related commands

Table 1. Command related to QUERY PROCESS

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
IDENTIFY DUPLICATES	Identifies duplicate data in a storage pool.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

QUERY PROFILE (Query a profile)

Use this command to display information about profiles and associated objects. Issue this command from a configuration manager or from a managed server. You can use this command to get profile information from any configuration manager defined to the server, even if the server does not subscribe to any profile.

If you query a locked profile from the configuration manager to which the profile belongs, complete profile information is displayed. If you query a locked profile from another server, the query displays only that the profile is locked.

Privilege class

Any administrator can issue this command.

Syntax

```

.*-----
>>-Query PROFILE--+----->
                    '-profile_name-'

>--+----->
|                               (1) |
'-SERVer-----server_name-----'

.-Format----Standard----- .-USELocal----Yes-----
>--+-----+-----><
'-Format-----Standard-+-' '-USELocal-----Yes-+-'
                    '-Detailed-'                    '-No--'

```

Notes:

1. The server name you specify depends on the server from which you issue the command. See the description of the SERVER parameter.

Parameters

profile_name

Specifies the profile to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all profiles.

SERVer

Specifies the configuration manager whose profile information is displayed. The requirements for the name depends on where the query is issued:

- From a configuration manager: This parameter is optional. The default is the configuration manager's name.
- From a managed server: This parameter is optional. The default is the name of the configuration manager for this managed server.
- From a server that is neither a configuration manager nor a managed server: You must specify a name.

Format

Specifies whether partial or detailed information is displayed. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that detailed information is displayed.

USELocal

When you perform the query from a managed server, this parameter specifies whether the profile information is obtained from the configuration manager or the managed server. If the profile information does not exist on the managed server, the information is obtained from the configuration manager, regardless of the value of this parameter.

If you use this parameter on a server that is not managed by the configuration manager that owns the profile, the parameter is ignored. The default value is YES. Possible values are:

Yes

Specifies that the profile information, if available, is obtained from the managed server. The configuration manager is contacted if information is not available from the managed server.

No

Specifies that the profile information is obtained from the configuration manager even if the information is available from the managed server. This ensures that you receive current information about the profile.

Example: List profiles from a configuration manager

Display profile information from a configuration manager. See Field descriptions for field descriptions.

```
query profile
```

Configuration manager	Profile name	Locked?
SERVER1	DEFAULT_PROFILE	No
SERVER1	ADMIN_INFO	No
SERVER1	EMPLOYEE	No
SERVER1	PERSONNEL	Yes

Example: Display detailed profile information for a managed server

From a managed server, display current detailed information for profile ADMIN_INFO. See Field descriptions for field descriptions. Note: When the profile is locked, most fields are not displayed.

```
query profile admin_info  
format=detailed uselocal=no
```

```
Configuration manager: SERVER1  
Profile name: ADMIN_INFO  
Locked: No  
Description: Distributed administrative schedules
```

```

Server administrators: DENNIS EMILY ANDREA
Policy domains: ADMIN RECORDS
Administrative command schedules: ** all objects **
Server Command Scripts:
Client Option Sets:
Servers:
Server Groups:

```

Field descriptions

Configuration manager
The name of the configuration manager that owns the profile.

Profile name
The name of the profile.

Locked?
Whether the profile is locked.

Description
The description of the profile.

Server administrators
The administrators that are associated with the profile.

Policy domains
The policy domains that are associated with the profile.

Administrative command schedules
The administrative schedules that are associated with the profile.

Server Command Scripts
The server command scripts that are associated with the profile.

Client Option Sets
The client option sets that are associated with the profile.

Servers
The servers that are associated with the profile.

Server Groups
The names of server groups that are associated with the profile.

Related commands

Table 1. Commands related to QUERY PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

QUERY PROTECTSTATUS (Query the status of storage pool protection)

Use this command to display information about the status of storage pool protection for directory-container storage pools.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*------.
>>-Query PROTECTStatus--+----->
      '-pool_name-'

.-Format----Standard----.
>--+-----+----->>
  '-Format----+Standard-+-'
      '-Detailed-'

```

Parameters

pool_name

Specifies the name of the directory-container storage pool to be queried. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value, the status of all directory-container storage pools is displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Specify one of the following values:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary information about a specific storage pool

Display information about the storage pool that is named POOL1. Issue the following command:

```
query protectstatus pool1
```

Source Server Name	Source Storage Pool	Target Server Name	Target Storage Pool	Pct. Protected	Last Complete Protect
NEXT	POOL1	NEXT	POOL1COPY	96.55	02/17/2017 11:15:07
NEXT	POOL1	NEXT1	POOL2	99.99	02/17/2017 11:14:53
NEXT	POOL1	UNKNOWN	UNKNOWN	UNKNOWN	02/17/2017 11:13:44
NEXT1	POOL2	NEXT	POOL1	100.00	02/17/2017 12:56:58

See Field descriptions for field descriptions.

Example: Display detailed information about a specific storage pool

Display information in full detail about the storage pool named, POOL1. Issue the following command:

```
query protectstatus pool1 format=detailed
```

```

  Source Server Name: NEXT
  Source Storage Pool: POOL1
  Target Server Name: NEXT
  Target Storage Pool: POOL1COPY
  Pct. Protected: 96.55
  Data Extents Protected: 1,747
  Data Extents Total: 1,852
  Protected (MB): 165.33
  Total (MB): 171.23
  Last Completed Protection: 02/17/2017 11:15:07
  Last Refresh Date/Time: 02/19/2017 00:27:12

```

See Field descriptions for field descriptions.

Field descriptions

Source Server Name

The name of the source server.

Source Storage Pool
The name of the directory-container storage pool on the source server.

Target Server Name
The name of the target server.

Target Storage Pool
The name of the directory-container storage pool on the target server.

Pct. Protected
The percentage of protected data in the directory-container storage pool.

Data Extents Protected
The number of data extents that are protected in the directory-container storage pool.

Data Extents Total
The total number of data extents in the directory-container storage pool.

Protected (MB)
The total amount of protected data that is in the directory-container storage pool, in megabytes.

Total (MB)
The total amount of data that is in the directory-container storage pool, in megabytes.

Last Completed Protection
The date and time that the directory-container storage pool was last protected.

Last Refresh Date/Time
The date and time that the directory-container storage pool was last refreshed.

Related commands

Table 1. Commands related to QUERY PROTECTSTATUS

Command	Description
PROTECT STGPOOL	Protects a directory-container storage pool.

QUERY PROXYNODE (Query proxy authority for a client node)

Use this command to display client nodes with authority to act as proxy to other client nodes in the IBM Spectrum Protect™ server.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query PROXynode----TArget-----.*-----
                                     +-----+----->>
                                     '-target_node_name-'
```

Parameters

TArget
Specifies the name of the node targeted by the node with proxy authority. It is optional to specify a target node name. Wildcard names can be used to specify the target node name. A comma-separated list of node names is also allowed.

Example: List client nodes with proxy authority

To display all IBM Spectrum Protect client nodes with proxy authority to the target node named MYCLUSTER, issue the following command.

```
query proxynode target=mycluster

Target Node      Agent Node
-----
FRED             MOE MINIE MICKEY
ALPHA           BETA GAMMA DELTA
```


Table 1. Sample output for several products managed by one IBM Spectrum Protect server

Product	Number of Client Devices	Number of Server Devices	PVU of Server Devices
IBM Spectrum Protect Extended Edition	1,000	905	90,500
IBM Spectrum Protect for Storage Area Networks	50	10	1,000
IBM Spectrum Protect for Space Management	0	0	0
IBM Spectrum Protect for Mail	0	25	5,000
IBM Spectrum Protect for Databases	0	1,025	20,500
IBM Spectrum Protect for Enterprise Resource Planning	0	25	5,000
IBM Spectrum Protect for System Backup and Recovery	0	0	0
Other Node Classifications	Number		
Nodes earlier than Version 6.3 with no PVU information available at this time	10		
Nodes at Version 6.3 or later with no PVU match	9		
Nodes classified by the administrator as "other-device"	8		
Nodes defined as a non-licensed API application	6		

The following list provides details about the example fields:

Product

The IBM Spectrum Protect product name.

Number of Client Devices

The estimated number of client devices that are managed by the product. By default, only nodes on Microsoft Windows 7, Microsoft Windows XP Professional, and Apple systems are assumed to be client devices.

Number of Server Devices

The estimated number of server devices that are managed by the product. By default, nodes on all platforms except for Microsoft Windows 7, Microsoft Windows XP Professional, and Apple systems are assumed to be server devices. This number also includes the server on which IBM Spectrum Protect is running.

PVU of Server Devices

The estimated PVUs of all nodes that are connected as server devices.

Nodes earlier than Version 6.3 with no PVU information available at this time

Devices that do not report processor information to the server.

Nodes at Version 6.3 or later with no PVU match

Devices that do not report all required values or some values were reported as "Unknown".

Nodes classified by the administrator as "other-device"

Nodes that are excluded from PVU counting by the administrator by using the update node roleoverride=other command.

Nodes defined as a non-licensed API application

Nodes such as DB2® backup or custom API applications.

Example: Display detailed node information

Display information for individual nodes by specifying the detailed (d) value for the Format parameter. Issue the following command:

```
tsm: PATMOS_630> query pvestimate f=d
```

Table 2. Node classifications for specific products

Product	Number of Client Devices	Number of Server Devices	PVU of Server Devices
IBM Spectrum Protect Extended Edition	1,000	905	90,500
- banode1	1		

Product	Number of Client Devices	Number of Server Devices	PVU of Server Devices
- banode2		1	200
- banode3	1		
- banode3		1	100
IBM Spectrum Protect for Storage Area Networks	50	10	1,000
- stagent1		1	50
- stagent2		1	100
IBM Spectrum Protect for Space Management	0	0	0
IBM Spectrum Protect for Mail	0	25	5,000
- mailnode1		1	200
- mailnode2		1	100
IBM Spectrum Protect for Databases	0	1,025	20,500
- dbnode1		1	200
- dbnode2		1	100
IBM Spectrum Protect for Enterprise Resource Planning	0	25	5,000
- erpnode1		1	50
- erpnode2		1	100
IBM Spectrum Protect for System Backup and Recovery	0	0	0
Other Node Classifications	Number		
Nodes earlier than Version 6.3 with no PVU information available at this time	10		
- oldnode1	1		
- oldnode2	1		
- mailnote44	1		
- erpnode66	1		
Nodes at Version 6.3 or later with no PVU match	10		
- badcitnode1	1		
- badcitnode2	1		
- mailnode23	1		
- erpnode34	1		
Nodes classified by administrator as "other-device"	8		
- overriddennode1	1		
- overriddennode2	1		

media_name

Specifies the name of the recovery media. You can use wildcard characters to specify the name. This parameter is optional. The default is all recovery media.

Type

Specifies the type of media to be queried. This parameter is optional. If this parameter is not specified, all recovery media are queried. Possible values are:

BOot

Only boot media are queried.

OTHer

All media other than boot media are queried.

LOcation

Specifies the location of the recovery media to be queried. This parameter is optional. You can specify up to 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Format

Specifies how the information is displayed. This parameter is optional. Possible values are:

Standard

Displays partial information. This is the default.

Detailed

Displays all information.

Example: Display summary information for a specific recovery media

Display information for the recovery media named RECMED1. See Field descriptions for field descriptions.

```
query recoverymedia RECMED1
```

Recovery Media Name	Volume Names	Location	Machine Name
RECMED1	vol1 vol2 vol3 vol4	IRONMOUNTAIN	MACH1

Example: Display detailed information for a specific recovery media

Display detailed information for the recovery media named RECMED1. See Field descriptions for field descriptions.

```
query recoverymedia RECMED1 format=detailed
```

```
Recovery Media Name: RECMED1
Type: Boot
Volume Names: vol1 vol2 vol3 vol4
Location: IRONMOUNTAIN
Description:
Product:
Product Information:
Machine Name: MACH1
```

Field descriptions

Recovery Media Name

The name of the recovery media.

Type

Whether the recovery media are boot media or another type of media. Possible values are:

Boot

The recovery media are boot media.

Other

The recovery media are not boot media.

Volume Names

The set of volumes that contain the data needed to recover machines associated with this media.

Location

Where the recovery media is stored.

Description

- A description of the recovery media.
- Product
 - The product used to create the boot media.
- Product Information
 - Information about the product that created the boot media. This information may be needed for restoring the machine.
- Machine Name
 - The machines that are associated with this recovery media.

Related commands

Table 1. Commands related to QUERY RECOVERYMEDIA

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

QUERY REPLICATION (Query node replication processes)

Use this command to display information about running and ended node-replication processes.

Issue this command on the server that acts as a source for replicated data.

Important: You cannot display information about running replication processes for client nodes that are being converted from import and export operations to replication operations. The conversion process might run for a long time, but it occurs only once for a client node that is being converted.

By default, records about completed node-replication processes are retained for 30 calendar days. A *calendar day* consists of 24-hours, from midnight to midnight.

To display the retention period, issue the QUERY STATUS command. Check the value in the Replication Record Retention Period field. To change the retention period, issue the SET REPLRETENTION command.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query REPLIcation--node_name----->
. -*-----
>--+-----+----->
| (1) |
'-----+--file_space_name--+--'
| '-----' |
'--+FSID--+-----'
| (2) |
|-----|

.-NAMEType---SERVER-----
>--+-----+----->
'-NAMEType---+SERVER---+
+UNICODE--+
| (2) |
'-FSID-----'

.-CODEType---BOTH-----
>--+-----+----->
'-CODEType---+BOTH---+
+UNICODE-----
'-NONUNICODE-'
```

```

.-DISplay---1-----
>-----+----->
'-DISplay---number_of_days-'

>-----+----->
'-PROcessid---process_identifier-'

.-Status---All----- .-Format---Standard-----
>-----+-----+----->>
'-Status---+All---+' '-Format---+Standard++'
          +-RUnning++          '-Detailed-'
          +-ENded---+
          '-FAiled--'

```

Notes:

1. Do not mix FSIDs (file space identifiers) and file space names in the same command.
2. Do not specify FSID if you use wildcard characters for the client node name.

Parameters

node_name (Required)

Specifies the name of the client node to be queried. You can use wildcard characters when you specify this name, with one exception. If the value of the NAMETYPE parameter is FSID, do not specify wildcard characters for the client node name. The FSID value indicates the file space identifier. File spaces with identical names can have different identifiers in different client nodes.

filesystem_name or FSID

Specifies the name of the file space or the file space identifier (FSID) to be queried. A name or FSID is optional. If you do not specify a name or an FSID, all file spaces are queried.

filesystem_name

Specifies the name of the file space that has data to be queried. File space names are case-sensitive. To determine the correct capitalization for the file space, issue the QUERY FILESPACE command. Separate multiple names with commas with no intervening spaces. When you specify a name, you can use wildcard characters.

A server that has clients with Unicode-enabled file spaces might have to convert the file space name. For example, the server might have to convert a name from the server code page to Unicode. For details, see the NAMETYPE parameter. If you do not specify a file space name, or if you specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

FSID

Specifies the file space identifier for the file space to be queried. The server uses FSIDs to find the file spaces to replicate. To determine the FSID for a file space, issue the QUERY FILESPACE command. Separate multiple FSIDs with commas with no intervening spaces. If you specify an FSID, the value of the NAMETYPE parameter must be FSID.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for IBM Spectrum Protect™ clients that are Unicode-enabled and that have Windows, Macintosh OS X, or NetWare operating systems.

Use this parameter only if you enter a partly qualified or fully qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret file space names.

UNICODE

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page. Conversion can also fail if the server cannot access system conversion routines.

FSID

The server interprets file space names by using their file space identifiers.

CODEType

Specifies the type of file spaces to be included in the query. The default value is BOTH, which means that file spaces are included regardless of code page type. Use this parameter only if you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include file spaces that are in Unicode only.

NONUNICODE

Include file spaces that are not in Unicode only.

BOTH

Include all file spaces regardless of code page type.

DISplay

Specifies the number of days of node replication history to display. The default value is 1, which displays information about running node replication processes and about processes that completed during the current calendar day. The maximum value is 9999.

You can specify a number that is the same as or less than the number of days that are specified as the retention period for the replication history records. If you specify a value that is more than the value of the replication retention period or more than the number of days that replication records are collected, the server displays only the number of replication history records that are available. For example, suppose that the replication retention period is 30 days and that the replication process is running for only 10 days. If you specify `DISPLAY=20`, only 10 days of replication history are displayed.

PROcessid

Specifies the node replication history that is associated with a particular process identified by the process identifier. This parameter is optional. If you do not specify this parameter, all processes are displayed for the number of days that are specified by the DISPLAY parameter.

Restarting the server can cause the server to reuse process IDs. Reuse of process IDs can result in duplicate process IDs for separate processes.

STatus

Specifies the status of the file spaces to query. This parameter is optional. The default value is ALL. You can specify one of the following values:

ALL

Specifies all file spaces that are replicating, file spaces that replicated successfully, and file spaces that did not finish replicating or replicated with errors.

RUNning

Specifies all file spaces that are replicating to the target replication server.

ENded

Specifies all file spaces that replicated successfully and file spaces that did not finish replicating or replicated with errors.

FAiled

Specifies all file spaces that did not finish replicating or replicated with errors.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. You can specify one of the following values:

Standard

Specifies that partial information is displayed for node replication processes.

Detailed

Specifies that all available information for the node replication processes is displayed.

Example: Display information about replication processes for a file space

Display information about replication processes for a file space in client node PAYROLL. The file space identifier is 10.

```
query replication ironman
```


NodeName	Filespace Name	FSID	Start Time	End Time	Status	Phase
IRONMAN	/space	2	02/08/11 21:44:19	02/08/11 21:48:14	Ended	None

query replication ironman format=detailed

```

Node Name: IRONMAN
Filespace Name: /space
FSID: 2
Start Time: 02/08/11 21:44:19
End Time: 02/08/11 21:48:14
Status: Ended
Process Number: 4
Command: replicate node ironman
Phase: None
Process Running Time: 0 Day(s) 0 Hour(s)
4 Minute(s)
Completion State: Complete
Reason For Incompletion: None
Backup Last Update Date/Time:
Backup Target Server:
Backup Files Needing No Action: 0
Backup Files To Replicate: 0
Backup Files Replicated: 0
Backup Files Not Replicated Due to Errors: 0
Backup Files Not Yet Replicated: 0
Backup Files To Delete: 0
Backup Files Deleted: 0
Backup Files Not Deleted Due To Errors: 0
Backup Files To Update: 0
Backup Files Updated: 0
Backup Files Not Updated Due To Errors: 0
Backup Bytes To Replicate (MB): 0
Backup Bytes Replicated (MB): 0
Backup Bytes Transferred (MB): 0
Backup Bytes Not Replicated Due
To Errors (MB): 0
Backup Bytes Not Yet Replicated (MB): 0

Archive Last Update Date/Time: 02/08/11 21:48:14
Archive Target Server: NIGLINA
Archive Files Needing No Action: 0
Archive Files To Replicate: 39,416
Archive Files Replicated: 39,206
Archive Files Not Replicated Due to Errors: 210
Archive Files Not Yet Replicated: 0
Archive Files To Delete: 0
Archive Files Deleted: 0
Archive Files Not Deleted Due To Errors: 0
Archive Files To Update: 0
Archive Files Updated: 0

Archive Files Not Updated Due To Errors: 0
Archive Bytes To Replicate (MB): 4,335
Archive Bytes Replicated (MB): 4,335
Archive Bytes Transferred (MB): 0
Archive Bytes Not Replicated
Due To Errors (MB): 0
Archive Bytes Not Yet Replicated (MB): 0

Space Managed Last Update Date/Time:
Space Management Target Server:
Space Managed Files Needing No Action: 0
Space Managed Files To Replicate: 0
Space Managed Files Replicated: 0
Space Managed Files Not Replicated
Due to Errors: 0
Space Managed Files Not Yet Replicated: 0
Space Managed Files To Delete: 0
Space Managed Files Deleted: 0
Space Managed Files Not Deleted
Due To Errors: 0

```

```

    Space Managed Files To Update: 0
    Space Managed Files Updated: 0
    Space Managed Files Not Updated
        Due To Errors: 0
    Space Managed Bytes To Replicate (MB): 0
    Space Managed Bytes Replicated (MB): 0
    Space Managed Bytes Transferred (MB): 0
    Space Managed Bytes Not Replicated
        Due To Errors (MB): 0
    Space Managed Bytes Not Yet Replicated (MB): 0
    Total Files Needing No Action: 0
    Total Files To Replicate: 39,416
    Total Files Replicated: 39,206
    Total Files Not Replicated Due To Errors: 210
    Total Files Not Yet Replicated: 0
    Total Files To Delete: 0
    Total Files Deleted: 0
    Total Files Not Deleted Due To Errors: 0
    Total Files To Update: 0
    Total Files Updated: 0
    Total Files Not Updated Due To Errors: 0
    Total Bytes To Replicate (MB): 4,335
    Total Bytes Replicated (MB): 4,335
    Total Bytes Transferred (MB):
    Total Bytes Not Replicated
        Due to Errors (MB):
    Total Bytes Not Yet Replicated (MB):
    Estimated Percentage Complete: 100
    Estimated Time Remaining:
    Estimated Time of Completion:

```

Field descriptions

Node Name

The name of the client node whose data is displayed.

Filespace Name

The name of the client file space whose data is displayed.

FSID

The file space identifier.

Start Time

The date and time that the node replication process started.

End Time

The date and time that the node replication process ended.

Status

The status of the node replication process. The following values are possible:

Running

The process is active and is either searching for eligible data or sending data to the target replication server.

Ended

The process ended or failed.

Failed

The process failed.

Process Number

The identifier for the node replication process.

The same process number can have different start times. If a replication process starts and the server is restarted, the server begins assigning process numbers that begin with the number 1. Replication processes that start after a server restart can obtain process numbers that are already assigned to other replication processes in the replication history. To identify unique replication processes, use the start time.

Command

The command that was issued to start the node replication process.

Phase

The phase of a running node-replication process. The following phases are listed in the order in which they occur:

Identifying

The node replication process started to identify data to be replicated, but data is not yet being sent to the target replication server.

Identifying and replicating

The node replication process is identifying data to be replicated and transferring the data to the target replication server.

Replicating

The node replication process identified the data and is transferring files to the target replication server.

None

The node replication process is not running.

Process Running Time

The running time of the node replication process.

Completion State

The state of the node replication process. The following values are possible:

Complete

The node replication process completed.

Incomplete

The node replication process ended without running to completion. To determine the reason, check the value in the Reason for Incompletion field.

Reason for Incompletion

The reason why the node replication process ended without completing. Possible values include *canceled* and *other*. The value *other* can indicate that the server was halted during replication or that the server failed.

Backup Last Update Date/Time

The date and time that statistics for backup were last updated. The specified time is the time that the files in the file space were identified for replication or when each batch of files was sent to the target replication server.

Archive Last Update Date/Time

The date and time that statistics for archive were last updated. The specified time is the time that the files in the file space were identified for replication or when each batch of files was sent to the target replication server.

Space Managed Last Update Date/Time

The date and time that statistics for space-managed files were last updated. The specified time is the time that the files in the file space were identified for replication or when each batch of files was sent to the target replication server.

Backup Target Server

The name of the target replication server for backup files.

Archive Target Server

The name of the target replication server for archive files.

Space Management Target Server

The name of the target replication server for space-managed files.

Backup Files Needing No Action

The number of backup files in the file space that did not need to be replicated, updated, or deleted.

Archive Files Needing No Action

The number of archive files in the file space that did not need to be replicated, updated, or deleted.

Space Managed Files Needing No Action

The number of space-managed files in the file space that did not need to be replicated, updated, or deleted.

Backup Files To Replicate

The number of backup files to replicate to the target replication server.

Archive Files To Replicate

The number of archive files to replicate to the target replication server.

Space Managed Files To Replicate

The number of space-managed files to replicate to the target replication server.

Backup Files Replicated

The number of backup files that are replicated to the target replication server.

Archive Files Replicated

The number of archive files that are replicated to the target replication server.

Space Managed Files Replicated

The number of space-managed files that are replicated to the target replication server.

Backup Files Not Replicated Due To Errors

The number of backup files that were not replicated to the target replication server because of errors.

Archive Files Not Replicated Due To Errors

The number of archive files that were not replicated to the target replication server because of errors.

Space Managed Files Not Replicated Due To Errors

The number of space-managed files that were not replicated to the target replication server because of errors.

Backup Files Not Yet Replicated

The number of backup files that are not yet replicated to the target replication server.

Archive Files Not Yet Replicated

The number of archive files that are not yet replicated to the target replication server.

Space Managed Files Not Yet Replicated

The number of space-managed files that are not yet replicated to the target replication server.

Backup Files To Delete

The number of backup files to be deleted on the target replication server.

Archive Files To Delete

The number of archive files to be deleted on the target replication server.

Space Managed Files To Delete

The number of space-managed files to be deleted on the target replication server.

Backup Files Deleted

The number of backup files that are deleted on the target replication server.

Archive Files Deleted

The number of archive files that are deleted on the target replication server.

Space Managed Files Deleted

The number of space-managed files that are deleted on the target replication server.

Backup Files Not Deleted Due To Errors

The number of backup files that were not deleted from the target replication server because of errors.

Archive Files Not Deleted Due To Errors

The number of archive files that were not deleted from the target replication server because of errors.

Space Managed Files Not Deleted Due To Errors

The number of space-managed files that were not deleted from the target replication server because of errors.

Backup Files To Update

The number of backup files to update on the target replication server. If the metadata of a file is changed, the changed fields are sent to the target replication server.

Archive Files To Update

The number of archive files to update on the target replication server. If the metadata of a file is changed, the changed fields are sent to the target replication server.

Space Managed Files To Update

The number of space-managed files to update on the target replication server. If the metadata of a file is changed, the changed fields are sent to the target replication server.

Backup Files Updated

The number of backup files that are updated on the target replication server.

Archive Files Updated

The number of archive files that are updated on the target replication server.

Space Managed Files Updated

The number of space-managed files that are updated on the target replication server.

Backup Files Not Updated Due To Errors

The number of backup files that were not updated on the target replication server because of errors.

Archive Files Not Updated Due To Errors

The number of archive files that were not updated on the target replication server because of errors.

Space Managed Files Not Updated Due To Errors

The number of space-managed files that were not updated on the target replication server because of errors.

Backup Bytes To Replicate (MB)

The number of backup bytes to replicate to the target replication server.

Archive Bytes To Replicate (MB)

The number of archive bytes to replicate to the target replication server.

Space Managed Bytes To Replicate (MB)

The number of space-managed bytes to replicate to the target replication server.

Backup Bytes Replicated (MB)

The number of backup bytes that are replicated to the target replication server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Archive Bytes Replicated (MB)

The number of archive bytes that are replicated to the target replication server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Space Managed Bytes Replicated (MB)

The number of space-managed bytes that are replicated to the target replication server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Backup Bytes Transferred (MB)

The number of backup bytes that were sent to the target replication server.

The value in this field represents the actual number of file bytes sent to the target replication server. This value is calculated by subtracting the number of bytes not sent because of deduplication from the number of bytes to replicate.

Archive Bytes Transferred (MB)

The number of archive bytes that were sent to the target replication server.

The value in this field represents the actual number of file bytes sent to the target replication server. This value is calculated by subtracting the number of bytes not sent because of deduplication from the number of bytes to replicate.

Space Managed Bytes Transferred (MB)

The number of space-managed bytes that were sent to the target replication server.

The value in this field represents the actual number of file bytes sent to the target replication server. This value is calculated by subtracting the number of bytes not sent because of deduplication from the number of bytes to replicate.

Backup Bytes Not Replicated Due to Errors (MB)

The number of backup bytes that were not replicated to the target replication server because of errors.

Archive Bytes Not Replicated Due to Errors (MB)

The number of archive bytes that were not replicated to the target replication server because of errors.

Space Managed Bytes Not Replicated Due to Errors (MB)

The number of space-managed bytes that were not replicated to the target replication server because of errors.

Backup Bytes Not Yet Replicated (MB)

The number of backup bytes not yet replicated to the target replication server.

Archive Bytes Not Yet Replicated (MB)

The number of archive bytes not yet replicated to the target replication server.

Space Managed Bytes Not Yet Replicated (MB)

The number of space-managed bytes not yet replicated to the target replication server.

Total Files Needing No Action

The total number of files in the file space that did not need to be replicated, updated, or deleted.

Total Files To Replicate

The total number of files to replicate to the target replication server.

Total Files Replicated

The total number of files that are replicated to the target replication server.

Total Files Not Replicated Due To Errors

The total number of files that were not replicated because of errors.

Total files Not Yet Replicated

The total number of files that are not yet replicated to the target replication server.

Total Files To Delete

The total number of files that were deleted on the target replication server.

Total Files Deleted

The total number of files that are deleted on the target replication server.

Total Files Not Deleted Due to Errors

The total number of backup, archive, and space-managed files that were not deleted on the target replication server because of errors.

Total Files To Update

The total number of files to be updated on the target replication server. When the metadata of a file is changed, the changed fields are sent to the target replication server.

Total Files Updated

The total number of files that are updated on the target replication server.

Total Files Not Updated Due to Errors

The total number of backup, archive, and space-managed files that were not updated on the target replication server because of errors.

Total Bytes To Replicate (MB)

The total number of bytes to replicate to the target replication server.

Total Bytes Replicated (MB)

The total number of bytes that are replicated to the target server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Total Bytes Transferred (MB)

The total number of bytes that were transferred to the target replication server.

For files stored in a deduplicated storage pool, the value in this field includes the number of bytes in the original file before duplicate extents were removed. If duplicate extents were already on the target replication server, the number of bytes in the original file is more than the number of bytes transferred.

Total Bytes Not Replicated Due to Errors (MB)

The total number of bytes that were skipped because the source replication server was unable to transfer them to the target replication server.

Total Bytes Not Yet Replicated (MB)

The total number of bytes not yet transferred to the target replication server.

Estimated Percentage Complete

The estimated completion percentage that is based on the number of bytes.

Estimated Time Remaining

The estimated time that remains before the node replication process is complete.

Estimated Time Of Completion

The estimated time when the node replication process ends.

Table 1. Commands related to QUERY REPLICATION

Command	Description
CANCEL REPLICATION	Cancels node replication processes.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PROCESS	Displays information about background processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY REPLRULE	Displays information about node replication rules.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET REPRETENTION	Specifies the retention period for replication history records.

QUERY REPLNODE (Display information about replication status for a client node)

Use this command to display the number of files that are stored for each replicated file space. Information is displayed about file spaces for every client node that is configured for replication.

A client node is configured for replication if it is enabled or disabled.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-.-.-.-.-
      v               |
>>-Query REPLNode-----node_name-----+----->
```

```
>-----<<
'-target_server_name-'
```

Parameters

node_name (Required)

Specifies the client node that owns the files about which you want information. You can specify one or more names. If you specify multiple names, separate the names with commas. Do not use intervening spaces. You can use wildcard characters to specify multiple names.

Information about client nodes that match the file criteria, but that are not configured for replication, is not displayed.

target_server_name

Specifies the name of the replication server to query for replication information. This parameter is optional. If you do not specify a value for this parameter, the server that is the default target for replicated data is queried.

As the value for this parameter, you can also specify a server that was formerly a target for replicated data.

The client nodes that are defined to a replication server can be the source or the target of replicated data. To determine whether a particular client node is sending or receiving data, issue the QUERY NODE command. Look for the value *Send* or *Receive* in the Replication Mode field of the output.

To display the name of the active target replication server, issue the QUERY STATUS command, and look for the name in the Target Replication Server field.

Example: List client node files on a source and a target replication server

The name of the client node is NODE1.

```
query replnode *
```

Node Name	Type	Filespace Name	FSID	Files on Server	Replication Server (1)	Files on Server (1)
NODE1	SpMg	/hmsmfs	1	1		
NODE1	Bkup	/lspace2	2	27		
NODE1	Arch	/lspace2	2	22	TGTSRV	22
NODE1	Bkup	/lspace	3	18,096		
NODE1	Arch	/lspace	3	61,150	TGTSRV	61,150
NODE2						

The number of files that are displayed for the replication servers might be different for the following reasons:

- The output of the QUERY REPLNODE command displays the number of files obtained from the occupancy table. The occupancy table contains only files that have a length greater than zero. Files that have a length of 0 and have been replicated are not reflected in this output.
- If only active data is replicated to the target server, the number of files that are displayed for the source server will be larger than the number of files that are displayed on the target server. The reason for the difference is that the source replication server has both active and inactive data, and the target server has only active data.
- A client node might have data that was exported from the source replication server and imported to the target replication server. If that data was synchronized and if the client node also stored data to the target replication server, then the number of files on the target replication server will be greater than the number of files stored as a result of export-and-import operations and replication.
- When you replicate node data from a source server prior to version 7.1, to a target server at version 7.1 or later, files that are larger than 10 GB are split in to smaller files if the SPLITLARGEOBJECTS parameter for the node definition is set to *Yes*. Each of these split files are counted on the target server.

Field descriptions

Node Name

The name of the client node that owns the files.

Type

The type of data. If this field is blank, the client node is configured for replication, but it does not have data on the replication server. In the example output, NODE2 is configured for replication, but it does not have backup, archive, or space-managed data.

The following values are possible:

Arch

Archive data

Bkup

Backup data

SpMg

Data that was migrated by IBM Spectrum Protect™ for Space Management clients

Filespace Name

The name of the file space that belongs to the node.

If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

FSID

The file space identifier for the file space. The server assigns a unique FSID when a file space is initially stored on the server. If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

Files on Server

The number of backup, archive, and space-managed files on the server on which this command is issued. If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

Replication Server (1)

The name of the replication server that is being queried for information. If this field is blank, one or more of the following conditions might exist:

- The file space of the node on the replication server where the command was issued does not have any data.
- The client node is not defined on replication server (1).
- The client node is defined on replication server (1), but the node is not configured for replication.
- The corresponding file space on replication server (1) does not have data or the file space is not defined.

Files on Server (1)

The number of files for the data type that are stored on the target replication server. This field can be blank. If it is, one or more of the following conditions might exist:

- Replication server (1) does not have any data.
- The client node is not defined on replication server (1).
- The client node is defined on replication server (1), but the node is not configured for replication.
- The corresponding file space on replication server (1) does not have data or the file space is not defined.

Related commands

Table 1. Commands related to QUERY REPLNODE

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.

Command	Description
QUERY REPLRULE	Displays information about node replication rules.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
UPDATE REPLRULE	Enables or disables replication rules.

QUERY REPLRULE (Query replication rules)

Use this command to display information about replication rules.

Issue this command on the server that acts as a source for replicated data.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query REPLRule .-*-----
                    +-----+
                    '|-+-ALL_DATA-----+|
                    |+-ACTIVE_DATA-----+
                    |+-ALL_DATA_HIGH_PRIORITY----+
                    '|-ACTIVE_DATA_HIGH_PRIORITY-|

```

Privilege class

Any administrator can issue this command.

Parameters

rule_name

Specifies the name of the replication rule that you want to display information about. This parameter is optional. You can use wildcard characters to specify one or more rules. If you do not specify this parameter, information about all rules is displayed in the query output. You can specify the following values:

ALL_DATA

Displays information about the ALL_DATA replication rule. This rule replicates backup, archive, or space-managed data. The data is replicated with a normal priority.

ACTIVE_DATA

Displays information about ACTIVE_DATA replication rule. This rule replicates only active backup data. The data is replicated with a normal priority. This rule is not valid for archive or space-managed data.

Attention: If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the `FORCERECONCILE=YES` parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Displays information about the ALL_DATA_HIGH_PRIORITY rule. This rule replicates backup, archive, or space-managed data. The data is replicated with a normal priority. In a replication process, high-priority data is replicated before normal-priority data.

ACTIVE_DATA_HIGH_PRIORITY

Displays information about the ACTIVE_DATA_HIGH_PRIORITY rule.

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority.

Example: Display information about a server replication rule

The name of the rule is ALL_DATA_HIGH_PRIORITY

```
query replrule all_data_high_priority
```

```
Replication Rule Name: ALL_DATA_HIGH_PRIORITY
Target Replication Server:
Active Only: No
Enabled: Yes
```

Field descriptions

Replication Rule Name

Specifies the name of the rule that was queried.

Target Replication Server

Specifies the name of the target replication server.

Active Only

Specifies whether the rule applies only to active backup data. The following values are possible:

Yes

Specifies that only active backup data is replicated for file spaces to which this rule is assigned.

No

Specifies that all backup data is replicated for file spaces to which this rule is assigned.

Enabled

Specifies whether the rule is enabled or disabled. The following values are possible:

Yes

Specifies that the rule is enabled for replication. Data in file spaces to which the rule is assigned is replicated.

No

Specifies that the rule is not enabled for replication. Data in file spaces to which the rule is assigned is not replicated.

Related commands

Table 1. Commands related to QUERY REPLRULE

Command	Description
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
UPDATE REPLRULE	Enables or disables replication rules.

QUERY REPLSERVER (Query a replication server)

Use this command to view information about all replication servers that are known server. The output from this command includes server information for the server from which the command was issued. The command indicates whether a replication server definition is deleted as a result of a REMOVE REPLSERVER command.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query REPLServer--+-----><
                    .-*-----
                    +-----+-----><
                    '-server_name-'
```

Example: Display summary statistics about all replicating servers

Display information about the replicating server. Issue the command from either the source or the target replication server:

```
query replserver *

Replication Globally Unique ID: 4d.83.fc.30.67.c1.11.e1.b8.
                                40.f0.de.f1.5e.f1.89
      Server Name: Server1
      Last Replication:
      Heartbeat:
Failover High Level Address: server1.example.com
Failover TCP Port Number: 1500
Failover SSL Port Number: 1542
      Deletion in Progress: No
      Dissimilar Policies:

Replication Globally Unique ID: 91.0f.ef.90.5c.cc.11.e1.ae.
                                34.08.00.27.00.58.dc
      Server Name: DRServer1
      Last Replication: 06/30/2012 08:16:30 PM
      Heartbeat: 07/09/2012 22:15:22 PM
Fail over High Level Address: drserver1.example.com
Failover TCP Port Number: 1500
Failover SSL Port Number: 1542
      Deletion in Progress: No
      Dissimilar Policies: On

Replication Globally Unique ID: 90.4f.53.b0.8e.cb.11.e3.a8.
                                2f.00.14.5e.55.b3.67
      Server Name: DRSERVER2
      Last Replication: 04/01/14 12:38:28
      Heartbeat: 05/29/14 11:15:44
Failover High Level Address: drserver2.example.com
Failover TCP Port Number: 1500
Failover SSL Port Number:
      Deletion in Progress: No
      Dissimilar Policies: Off
```

Example: Display summary statistics about a specific replicating server

Display information about the replicating server DRServer1. Issue the command from either the source or the target replication server:

```
query replserver drserver1

Replication Globally Unique ID: 91.0f.ef.90.5c.cc.11.e1.ae.
                                34.08.00.27.00.58.dc
      Server Name: DRServer1
      Last Replication: 06/30/2012 08:16:30 PM
      Heartbeat: 07/09/2012 22:15:22 PM
Fail over High Level Address: drserver1.example.com
Failover TCP Port Number: 1500
Failover SSL Port Number: 1542
      Deletion in Progress: No
      Dissimilar Policies: On
```

Parameters

server_name

Specifies the name of the replication server to be queried. You can use wildcard characters to specify this name. All matching servers are queried. If you do not specify a value for this parameter, all servers are queried. The parameter is optional.

Field descriptions

Replication Globally Unique ID

The unique identifier for the IBM Spectrum Protect™ server. The values for the Replication Globally Unique ID are created when a server is first used in a replication process.

Tip: The ID listed in the Replication Globally Unique ID field is not the same value as the value for the ID listed in the Machine Globally Unique ID field that is shown in the QUERY STATUS command.

Server Name

- The name of the replication server.
- Last Replication
 - The date of the last replication process that used the server.
- Heartbeat
 - The last time that the server completed a successful test communication session.
- Failover TCP Port Number
 - The active Transmission Control Protocol (TCP) client port on the replication server that is used for client connections. If the client is configured for TCP, the port is used to connect to the failover server.
- Failover SSL Port Number
 - The active Secure Sockets Layer (SSL) port on the replication server that is used for client connections. If the client is configured for SSL, the port is used to connect to the failover server.
- Failover High Level Address
 - The high-level address that the client uses to connect to the replication server during failover.
- Deletion in Progress
 - Specifies whether a REMOVE REPLSERVER command was issued for this replication server and is still in progress. The following values are possible:
 - Yes
 - The deletion of the replication server is in progress.
 - No
 - The deletion of the replication server is not in progress.
- Dissimilar Policies
 - Specifies whether the policies that are defined on the target replication server are enabled. The following values are possible:
 - On
 - The policies on the target replication server manage replicated client-node data.
 - Off
 - The policies on the source replication server manage replicated client-node data.

Related commands

Table 1. Commands related to QUERY REPLSERVER

Command	Description
REMOVE REPLNODE (Remove a client node from replication)	Removes a node from replication.
REMOVE REPLSERVER (Remove a replication server)	Removes a server from replication.

QUERY REQUEST (Query one or more pending mount requests)

Use the QUERY REQUEST command to show information about one or more pending mount requests. The server makes requests for the administrator to complete an action, like inserting a tape volume in a library after a CHECKIN LIBVOL is issued.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query REQuest-+-----+-----><
                    '-request_number-'
```

Parameters

request_number
 Specifies the identification number of the pending mount request. This parameter is optional. The default is all pending mount requests.

Example: List all pending mount requests

Display information about all pending mount requests after a CHECKIN LIBVOL is issued.

query request

Output for a manual Library

AIX

```
ANR8352I Requests outstanding:
ANR8326I 001: Mount 8MM volume EXP001 R/W
in drive 8MM.1 (/dev/mt0) of library
MANUALLIB within 60 minute(s).
```

Linux

```
ANR8352I Requests outstanding:
ANR8326I 001: Mount 8MM volume EXP001 R/W
in drive 8MM.1 (/dev/mt0) of library
MANUALLIB within 60 minute(s).
```

Windows

```
ANR8352I Requests outstanding:
ANR8326I 001: Mount GENERICTAPE volume EXP001 R/W
in drive 8MM.1 (mt3.0.0.0) of library
MANUALLIB within 60 minute(s).
```

Output for an automated Library

AIX

Windows

```
ANR8352I Requests outstanding:
ANR8306I 001: Insert LTO volume 133540L5 R/W into the slot with
element number 31 of library LTOLIB within 60 minutes; issue
'REPLY' along with the request ID when ready.
```

Linux

```
ANR8352I Requests outstanding:
ANR8306I 001: Insert 3590 volume 133540 R/W into the slot with element
number 31 of library 3590LIB within 60 minutes; issue 'REPLY'
along with the request ID when ready.
```

Related commands

Table 1. Related commands for QUERY REQUEST

Command	Description
CANCEL REQUEST	Cancels pending volume mount requests.
REPLY	Allows a request to continue processing.

QUERY RESTORE (Query restartable restore sessions)

Use this command to display information about the restartable restore sessions.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query--REStore--+-+-----+-----+-----+----->
                    '-node_name-' '-file_space_name-'
                    .-Format-----Standard----- .-NAMEType-----SERVER-----.
```

```
>-----+-----+-----+-----><
'-Format-----+Standard-+-' '-NAMEType-----+SERVER---+-'
          '-Detailed-'                +-UNICODE+-
          '-FSID-----'
```

Parameters

node_name

Specifies the client node to be queried. This parameter is optional. If you do not specify a value, all client nodes with restartable restore sessions are displayed. You must specify a value for this parameter if you specify a file space name.

file_space_name

Specifies the file space to be queried. This parameter is optional. If you do not specify a value, all file spaces are matched for the specified node.

For a server that has clients with support for Unicode, you may need to have the server convert the file space name that you enter. For example, you may need to have the server convert the name you enter from the server's code page to Unicode. See the NAMEType parameter for details.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect™ clients using Windows, Macintosh OS 9, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

Example: Display a restartable restore session on a specific client node

Display detailed information about client node JAMES associated with file space DRIVE_F_R. See Field descriptions for field descriptions.

```
query restore james drive_f_r format=detailed
```

```
Sess Number: -1
Restore State: Restartable
Elapsed Minutes: 2
Node Name: JAMES
FSID: 1
Filespace Name: DRIVE_F_R:
File Spec: /RESTORE/TESTDIR\
```

Field descriptions

Sess Number

Specifies the session number for the restartable restore session. The number for active restore sessions is the same number displayed on the QUERY SESSION command. For restore sessions in the restartable state, a negative number is

displayed for the session number. Any session number displayed in the QUERY RESTORE output may be specified from the QUERY RESTORE output.

Restore State

- Active: Specifies the restore session is actively restoring files to the client.
- Restartable: Specifies the restore session failed and can be restarted from where it left off.

Elapsed Minutes

Specifies the number of minutes since the restore session started. Any restartable restore session with an elapsed time greater than the RESTOREINTERVAL server option can be automatically deleted from the database when needed or during expiration processing. If the elapsed time is less than the RESTOREINTERVAL, you can delete this entry (and unlock the file space) only by issuing the CANCEL RESTORE command lowering the RESTOREINTERVAL value.

Node Name

Specifies the node associated with the restartable restore session.

FSID

Specifies the file space ID of the file space.

Filespace Name

Specifies the file space associated with the restartable restore session.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

File Spec

Specifies the file specification used on the restore operation. The same file specification must be specified if a failed restore operation is to be restarted from where it stopped.

Related commands

Table 1. Commands related to QUERY RESTORE

Command	Description
CANCEL RESTORE	Cancels a restartable restore session.

QUERY RPFCONTENT (Query recovery plan file contents stored on a target server)

Use this command to display the contents of a recovery plan file stored on a target server (that is, when the DEVCLASS parameter was specified on the PREPARE command). You can issue this command from either the server that created the file (the source server) or the server that stores the recovery plan file (the target server). You cannot issue this command from the server console.

The output may be delayed if the file is on tape.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Query RPFContent--plan_file_name----->
>--+DEVclass--==--device_class_name-+-----<
  '-NODEName--==--node_name-----'
```

Parameters

plan_file_name (Required)

Specifies the name of the recovery plan file to be queried. The format of the file name is servername.yyyymmdd.hhmmss. To see the names of existing files, issue the QUERY RPFFILE command.

DEVclass

Specifies the name of the device class used to create the recovery plan file. Wildcard characters are not allowed. Specify this parameter when:

- You want to display the contents of the recovery plan file that was created for this server.
- You are issuing this command to the same server on which the PREPARE command was issued (the source server).
- The specified device class name was used on the PREPARE command that created the recovery plan file.

NODENAME

Specifies the node name, registered on the target server, of the source server that created the recovery plan file. Wildcard characters are not allowed.

Specify this parameter when:

- You want to display the contents of the recovery plan file that was stored on this server.
- You are issuing this command to the server that was the target of the PREPARE command that created the recovery plan file.
- The specified node name is registered to this server with a node type of SERVER.
- The IBM Spectrum Protect™ server that created the recovery plan file is not available.

Example: Display the source server recovery plan

On the source server, display the contents of a recovery plan file that was created for this server on March 19, 1998, at 6:10 A.M. The PREPARE command specifies the device class REMOTE. The output of this command is the entire contents of the recovery plan file.

```
query rpfcontent branch1.19980319.061000 devclass=remote
```

Example: Display the target server recovery plan

On the target server, display the contents of a recovery plan file that was stored in this server on March 19, 1998, at 6:10 A.M. The server that created the file is registered on the target server as a node named POLARIS with a node type of SERVER. The output of this command is the entire contents of the recovery plan file.

```
query rpfcontent branch1.19980319.061000 nodename=polaris
```

Related commands

Table 1. Commands related to QUERY RPFCONTENT

Command	Description
PREPARE	Creates a recovery plan file.
QUERY RPFFILE	Displays information about recovery plan files.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.

Related information:

[Disaster recovery plan file](#)

QUERY RPFFILE (Query recovery plan file information stored on a target server)

Use this command to display information about recovery plan files stored on a target server. You can issue this command from either the server that created the file (the source server) or the server that stores the recovery plan file (the target server).

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query RPFile---+DEVclass-----device_class_name+----->
                '-NODENAME-----node_name-----'

    .-Source-----DBBackup----- .-Format-----Standard-----
>--+-----+-----+-----+-----+-----+-----+----->>
    '-Source-----+DBBackup-----' '-Format-----+Standard+-'
                '-DBSnapshot-'                '-Detailed-'
```

Parameters

DEVclass

Specifies the name of the device class that was used to create the recovery plan files. Use this parameter when logged on to the server that created the recovery plan file. You can use wildcard characters in the device class name. All recovery plan files that are created with the device class specified are included in the query.

NODENAME

Specifies the node name, registered on the target server, of the source server that created the recovery plan files. Use this parameter when logged on to the target server. You can use this parameter when the source server is not available. You can use wildcard characters to specify the node name. All file objects that are stored with the node name specified are included in this query.

Source

Specifies the type of database backup that was specified when the recovery plan file was prepared. This parameter is optional. The default is DBBACKUP. Possible values are:

DBBackup

The recovery plan file was prepared with full and incremental database backups specified.

DBSnapshot

The recovery plan file was prepared with snapshot database backups specified.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Displays partial information for the recovery plan file.

Detailed

Displays all information for the recovery plan file.

Example: Display detailed information about the recovery plans

Display recovery plan files that were created for this server using the specified device class. See Field descriptions for field descriptions.

```
query rpf file devclass=* format=detailed

Recovery Plan File Name: ALASKA.20000406.170423
      Node Name: BRANCH1
      Device Class Name: REMOTE
Recovery Plan File Type: RPF FILE
      Mgmt Class Name: STANDARD
Recovery Plan File Size: 16,255 Bytes
      Marked for Deletion: Yes
      Deletion Date: 06/12/2000 13:05:31

Recovery Plan File Name: ALASKA.20000407.170845
      Node Name: BRANCH1
      Device Class Name: REMOTE
Recovery Plan File Type: RPF SNAPSHOT
      Mgmt Class Name: STANDARD
Recovery Plan File Size: 16,425 Bytes
      Marked for Deletion: No
      Deletion Date:
```

Example: Display a list of recovery plans for a specific node name

Display a list of all recovery plan file objects that are stored with the specified node name (TYPE=SERVER). See Field descriptions for field descriptions.

```
query rprofile nodename=branch1
```

Recovery Plan File Name	Node Name	Device Class Name
ALASKA.19980406.170423	BRANCH1	REMOTE
ALASKA.19980407.170845	BRANCH1	REMOTE

Field descriptions

Recovery Plan File Name

The recovery plan file name.

Node Name

The node name that is registered with the target server and used to store the recovery plan file objects.

Device Class Name

The device class name that is defined in the source server and used to create the recovery plan files.

Recovery Plan File Type

The type of recovery plan file:

RPFIL

The plan assumes full plus incremental database backups.

RPFNSNAPSHOT

The plan assumes snapshot database backups.

Mgmt Class Name

The management class name that the recovery plan file is associated with in the target server.

Recovery Plan File Size

Estimated size of the recovery plan file object on the target server.

Marked For Deletion

Whether the object that contains the recovery plan file has been deleted from the source server and marked for deletion on the target server if the grace period has not expired. Possible values are:

Yes

The object is marked for deletion.

No

The object is not marked for deletion.

Deletion Date

Date that the object has been deleted from the source server and marked for deletion on the target server. This field is blank if the object has not been marked for deletion.

Related commands

Table 1. Commands related to QUERY RPFIL

Command	Description
PREPARE	Creates a recovery plan file.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.

QUERY SAN (Query the devices on the SAN)

Use this command to obtain information about devices that can be detected on a storage area network (SAN) so that you can configure IBM Spectrum Protect™ for LAN-free data movement.

AIX The QUERY SAN command requires the libhbaapi.a that supports SNIA common Host Bus Adapter (HBA) API. With this library object, IBM Spectrum Protect can call the hbaapi functions that are specified in the SNIA common HBA API standard.

Windows The QUERY SAN command requires the hbaapi.dll that supports SNIA common Host Bus Adapter (HBA) API. With this library object, IBM Spectrum Protect can call the hbaapi functions that are specified in the SNIA common HBAAPI standard.

Linux The QUERY SAN command requires the libhbaapi.so that supports SNIA common Host Bus Adapter (HBA) API. With this library object, IBM Spectrum Protect can call the hbaapi functions that are specified in the SNIA common HBAAPI standard. The QUERY SAN command might not show all the devices if the SANDISCOVERY server option is not set to ON.

Privilege class

Any administrator can issue this command.

Syntax

```

.-Type-----Any-----
>>-Query SAN-----+----->
      '-Type-----+Any-----+'
              +-Drive---+
              '-LIBRARY-'

.-Format-----Standard-----
>--+-----+----->>
      '-Format-----+Standard+-'
              '-Detailed-'

```

Parameters

Type

Specifies the type of device that is displayed. This parameter is optional. The default value is Any. Possible values are:

Any

Specifies that any device detected on the SAN is displayed.

DRive

Specifies that only drive devices are displayed.

LIBRARY

Specifies that only library devices are displayed.

Format

Specifies the type of information that is displayed. This parameter is optional. The default value is Standard. Possible values are:

Standard

Specifies that the information displayed is summarized.

Detailed

Specifies that complete information is displayed.

Tip: The output might not display the serial number of the device. If this happens, look on the back of the device or contact the manufacturer of the device.

Example: List drive devices

Display summary information for drive devices on a SAN. See Field descriptions for field descriptions.

```
query san type=drive
```

Device Type	Vendor	Product	Serial	Device
LIBRARY	STK	L180	MPC01000128	/dev/smc1
DRIVE	STK	9840D	331001017229	/dev/rmt3
DRIVE	Quantum	DLT4000	JF62806275	/dev/rmt4
DRIVE	Quantum	DLT4000	JP73213185	/dev/rmt5
DRIVE	STK	9840D	331000028779	/dev/rmt6

Example: Display drive device information

Display detailed information for all drive devices on a SAN. See Field descriptions for field descriptions.

```
query san type=drive format=detailed
```

```
Device Type:  DRIVE
Vendor:       IBM
Product:     03570B02
Serial Number:
Device:      mt10.2.0.3
DataMover:   No
Node WWN:    5005076206039E05
Port WWN:    5005076206439E05
LUN:         0
SCSI Port:   3
SCSI Bus:    0
SCSI Target: 10
```

Field descriptions

Device Type

The type of device that is being displayed.

Vendor

The name of the device's vendor.

Product

The name of the product that is assigned by the vendor.

Serial Number

The serial number of the device.

Device

The device special file name.

Data Mover

Whether the device is a data mover.

Node WWN

The worldwide name for the device.

Port WWN

The worldwide name for the device, which is specific to the port that the device is connected to.

LUN

The Logical Unit Number of the device.

SCSI Port

The port of the Fibre Channel (or SCSI) Host Bus Adapter.

SCSI Bus

The bus of the Host Bus Adapter card.

SCSI Target

The target number of the device.

Related commands

Table 1. Commands related to QUERY SAN

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.

QUERY SCHEDULE (Query schedules)

Use this command to display information about one or more schedules.

The QUERY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each operation are defined separately. Some options in the query display will be blank depending on whether the schedule style is classic or enhanced.

Table 1. Commands related to QUERY SCHEDULE

Command	Description
---------	-------------

The standard format displays a blank in the period column and an asterisk in the day column for enhanced schedules. To display complete information about an enhanced schedule, issue `FORMAT=DETAILED`.

Example: List schedules for a specific policy domain

Display all schedules that belong to the `EMPLOYEE_RECORDS` policy domain. See Field descriptions: Schedules for a specific policy domain for field descriptions.

```
query schedule employee_records
```

The standard format displays a blank in the period column and an asterisk in the day column for enhanced schedules. To display complete information about an enhanced schedule, issue `FORMAT=DETAILED`.

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
EMPLOYEE_RECORDS	WEEKLY_BACKUP	Inc Bk	2004.06.04 17.04.20	1 H	1 D	Any
EMPLOYEE_RECORDS	EMPLOYEE_BACKUP	Inc Bk	2004.06.04 17.04.20	1 H		(*)

Field descriptions: Schedules for a specific policy domain

Domain

Specifies the name of the policy domain to which the specified schedule belongs.

* (Asterisk)

Specifies whether the corresponding schedule has expired. If there is an asterisk in this column, the corresponding schedule has expired.

Schedule Name

Specifies the name of the schedule.

Action

Specifies the action that occurs when this schedule is processed.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window for this schedule.

Period

Specifies the time between startup windows (assuming `DAYOFWEEK=ANY`). The column is blank for enhanced schedules.

Day

Specifies the day of the week on which the startup windows for the schedule begin. The column contains an asterisk for enhanced schedules.

Example: Display detailed client schedules

From a managed server, display detailed information about client schedules. See Field descriptions: Detailed client schedules for field descriptions.

```
query schedule * type=client format=detailed
```

```

Policy Domain Name: ADMIN_RECORDS
Schedule Name: ADMIN_BACKUP
Description:
  Action: Backup
  Subaction: vApp
  Options:
  Objects:
  Priority: 5
Start Date/Time: 04/06/2013 17.04.20
Duration: 1 Hour(s)
Maximum Run Time (Minutes): 0
Schedule Style: Classic
  Period: 1 Day(s)
  Day of Week: Any
  Month:
  Day of Month:

```

```

Week of Month:
Expiration:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 04/06/2013 17.51.49
Managing profile: ADMIN_INFO

Policy Domain Name: EMPLOYEE_RECORDS
Schedule Name: EMPLOYEE_BACKUP
Description:
Action: Incremental
Subaction:
Options:
Objects:
Priority: 5
Start Date/Time: 2004.06.04 17.04.33
Duration: 1 Hour(s)
Maximum Run Time (Minutes): 0
Schedule Style: Enhanced
Period:
Day of Week: Any
Month: Mar, Jun, Nov
Day of Month: -14,14,22
Week of Month: Last
Expiration:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2004.06.04 17.18.30
Managing profile: EMPLOYEE

```

Field descriptions: Detailed client schedules

Policy Domain Name

Specifies the name of the policy domain.

Schedule Name

Specifies the name of the schedule.

Description

Specifies the description of the schedule.

Action

Specifies the type of action that occurs when this schedule is run. See the DEFINE SCHEDULE command for a listing of actions.

Subaction

Specifies that the type of operation identified by the ACTION parameter is to be scheduled. See the DEFINE SCHEDULE command for a listing of subactions.

Options

Specifies the options that are supplied to the DSMC command when the schedule is run.

Objects

Specifies the objects for which the specified action is performed.

Priority

Specifies the priority value for the schedule.

Start Date/Time

Specifies the initial starting date and time for the schedule.

Duration

Specifies the length of the startup window for the schedule.

Maximum Run Time (Minutes)

Specifies the number of minutes during which all client sessions that are started by the scheduled operation should be completed. If sessions are still running after the maximum run time, the server issues a warning message, but the sessions continue to run.

Schedule Style

Specifies whether classic or enhanced schedule rules are used.

Period

Specifies the time between startup windows (assuming DAYOFWEEK=ANY). This is not displayed for enhanced syntax schedules.

Day of Week

Specifies the day of the week on which the startup windows for the schedule begin. Using a standard format displays an asterisk in the day of week field for enhanced schedules.

Month

Specifies the months during which the schedule will run. This is not displayed for classic syntax schedules.

Day of Month

Specifies the days of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Week of Month

Specifies the weeks (first, second, third, fourth, or last) of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Expiration

Specifies the date and time on which this schedule expires. If this column is blank, the schedule does not expire.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the schedule. If this field contains a \$\$CONFIG_MANAGER\$\$, the schedule is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

Specifies the last date and time the schedule was last updated.

Managing Profile

Specifies the profile or profiles to which the managed server subscribed to get the definition of this schedule.

QUERY SCHEDULE (Query an administrative schedule)

Use this command to display information about one or more administrative schedules.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*-----
>>-Query SCHEDULE-----+-----+---Type-----Administrative--->
      '-schedule_name-'

      .-Format-----Standard-----
>--+-----+-----+-----+----->>
      '-Format-----+Standard+-'
      '-Detailed-'
```

Parameters

schedule_name

Specifies the name of the schedule to be queried. You can use a wildcard character to specify this name.

Type=Administrative (Required)

Specifies that the query displays administrative command schedules.

Format

Specifies how information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the schedules.

Detailed

Specifies that detailed information is displayed for the schedules.

The standard format displays a blank period column and an asterisk in the day column for enhanced schedules. Issue FORMAT=DETAILED to display complete information about an enhanced schedule.

Example: Display detailed information on administrative command schedules

From a managed server, display detailed information about administrative command schedules. See Field descriptions for field descriptions.

```
query schedule * type=administrative
format=detailed
```

```
Schedule Name: BACKUP_ARCHIVEPOOL
Description:
Command: backup db
```



```

        Priority: 5
        Start Date/Time: 2004.06.04 16.57.15
        Duration: 1 Hour(s)
Maximum Run Time (Minutes): 0
        Schedule Style: Classic
        Period: 1 Day(s)
        Day of Week: Any
        Month:
        Day of Month:
        Week of Month:
        Expiration:
        Active: No
Last Update by (administrator): $$CONFIG MANAGER$$
        Last Update Date/Time: 2004.06.04 17.51.49
        Managing Profile: ADMIN_INFO

        Schedule Name: MONTHLY_BACKUP
        Description:
        Command: q status
        Priority: 5
        Start Date/Time: 2004.06.04 16.57.14
        Duration: 1 Hour(s)
Maximum Run Time (Minutes): 0
        Schedule Style: Enhanced
        Period:
        Day of Week: Tue,Thu,Fri
        Month: Aug,Nov
        Day of Month:
        Week of Month: Second,Third
        Expiration:
        Active: No
Last Update by (administrator): $$CONFIG MANAGER
        Last Update Date/Time: 2004.06.04 17.51.49
        Managing Profile: ADMIN_INFO

```

Field descriptions

Schedule Name

Specifies the name of the schedule.

Description

Specifies the description of the schedule.

Command

Specifies the command that is scheduled.

Priority

Specifies the priority value for this schedule.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window.

Maximum Run Time (Minutes)

Specifies the number of minutes during which server processes that are started by the scheduled commands must be completed. If processes are still running after the maximum run time, the central scheduler cancels the processes.

Tips:

- This parameter does not apply to some processes, such as duplicate-identification processes, which can continue to run after the maximum run time.
- Another cancel time might be associated with some commands. For example, the MIGRATE STGPOOL command can include a parameter that specifies the length of time that the storage pool migration runs before the migration is automatically canceled. If you schedule a command for which a cancel time is defined, and you also define a maximum run time for the schedule, the processes are canceled at whichever cancel time is reached first.

Schedule Style

Specifies whether classic or enhanced schedule rules are used.

Period

Specifies the time between startup windows (assuming DAYOFWEEK=ANY). This is not displayed for enhanced syntax schedules.

Day of Week

Query a database that stores information about the location of all administrators.

```
query scratchpadentry admin_info location
```

```
Scratchpad major category: admin_info
  Scratchpad minor category: location
    Scratchpad subject: codjo
    Scratchpad line number: 1
      Scratchpad data: Toronto 5A24
      Date/time of creation: 2013-09-10, 10:15:50
      Last Update Date/Time: 2013-09-10, 10:15:50
Last Update by (administrator): CODJO
```

```
Scratchpad major category: admin_info
  Scratchpad minor category: location
    Scratchpad subject: jane
    Scratchpad line number: 1
      Scratchpad data: Raleigh GF85
      Date/time of creation: 2013-09-09, 14:29:40
      Last Update Date/Time: 2013-09-09, 14:29:40
Last Update by (administrator): JANE_W
```

```
Scratchpad major category: admin_info
  Scratchpad minor category: location
    Scratchpad subject: jane
    Scratchpad line number: 2
      Scratchpad data: Out of the office from 1-15 Nov.
      Date/time of creation: 2013-09-09, 14:30:05
      Last Update Date/Time: 2013-10-31, 16:55:52
Last Update by (administrator): JANE_W
```

```
Scratchpad major category: admin_info
  Scratchpad minor category: location
    Scratchpad subject: montse
    Scratchpad line number: 1
      Scratchpad data: Barcelona B19
      Date/time of creation: 2013-09-10, 04:34:37
      Last Update Date/Time: 2013-09-10, 04:34:37
Last Update by (administrator): MONTSERRAT
```

Field descriptions

Scratchpad data

The data that is stored in the scratch pad entry.

Date/time of creation

The date and time at which the scratch pad entry was created.

Last Update Date/Time

The date and time at which the scratch pad entry was last updated.

Last Update by (administrator)

The administrator who last updated the scratch pad entry.

Related commands

Table 1. Commands related to QUERY SCRATCHPADENTRY

Command	Description
DEFINE SCRATCHPADENTRY	Creates a line of data in the scratch pad.
DELETE SCRATCHPADENTRY	Deletes a line of data from the scratch pad.
SET SCRATCHPADRETENTION	Specifies the amount of time for which scratch pad entries are retained.
UPDATE SCRATCHPADENTRY	Updates data on a line in the scratch pad.

QUERY SCRIPT (Query IBM Spectrum Protect scripts)

Use this command to display information about scripts.

QCOLS	Display columns for a specified SQL table
QSAMPLE	Sample SQL Query
EXAMPLE	Backup the store pools and database when no sessions

Example: Display the contents of a script with line numbers

Display the lines of information for a script named Q_AUTHORITY.

```
query script q_authority format=lines
```

Name	Line Number	Command
Q_AUTHORITY	1	/* -----*/
	5	/* Script Name: Q_AUTHORITY */
	10	/* Description: Display administrators that */
	15	/* have the authority to issue */
	20	/* commands requiring a */
	25	/* specific privilege. */
	30	/* Parameter 1: privilege name - in the form */
	35	/* x_priv - EX. policy_priv */
	40	/* Example: run q_authority storage_priv */
	45	/* -----*/
	50	select admin_name from admins where -
	55	upper(system_priv) <> 'NO' or -
	60	upper(\$1) <> 'NO'

Example: Create a script from an existing script

Query the ENGDEV script and direct the output to a file named MY.SCRIPT.

```
query script engdev format=raw outputfile=my.script
```

Example: Display detailed script information

Display detailed information about scripts. See Field descriptions for field descriptions.

```
query script * format=detailed
```

Name:	QCOLS
Line Number:	DESCRIPTION
Command:	Display columns for a specified SQL table
Last Update by (administrator):	SERVER_CONSOLE
Last Update Date/Time:	12/02/1997 16:05:29
Name:	QCOLS
Line Number:	1
Command:	select colname from columns where tabname='\$1'
Last Update by (administrator):	SERVER_CONSOLE
Last Update Date/Time:	12/02/1997 16:05:29

Field descriptions

Name

The name of the script.

Line Number

The line number of the script or the string DESCRIPTION.

Command

The command included on the line number that is displayed in the previous field.

Last Update by (administrator)

The name of the administrator that defined or most recently updated the script.

Last Update Date/Time

The date and time that the administrator defined or updated the script.

Related commands

Table 1. Commands related to QUERY SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

Related concepts:

Server scripts

QUERY SERVER (Query a server)

Use this command to display information about a server definition.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query SERver .-*----- .-Format----Standard-----
                +-----+-----+-----+----->>
                '-server_name-' '-Format----+Standard-+-'
                                     '-Detailed-'
```

Parameters

server_name

Specifies the name of the server to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is all server names.

Format

Specifies how the information is displayed. The parameter is optional. The default is STANDARD.

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List all servers

Display information in standard format about all servers. See Field descriptions for field descriptions.

```
query server *
```

Server Name	Comm. Method	High-level Address	Low-level Address	Days Since Last Access	Server Password Set	Virtual Volume Password Set	Allow Replacement
SERVER_A	TCPIP	9.115.35.6	1501	11	Yes	No	No
SERVER_B	TCPIP	9.115.45.24	1500	<1	Yes	No	No
ASTRO	TCPIP	9.115.32.21	1500	24	Yes	No	No

Example: Display detailed information about a specific server

From a managed server, display detailed information about SERVER_A. See Field descriptions for field descriptions.

query server server_a format=detailed

```
Server Name: SERVER_A
Comm. Method: TCPIP
Transfer Method: TCPIP
High-level Address: 9.115.4.15
Low-level Address: 1500
Description:
Allow Replacement: No
Node Name:
Last Access Date/Time: 07/09/2013 09:00:00
Days Since Last Access: <1
Compression: Client's choice
Archive Delete Allowed?: No
URL:
Registration Date/Time: 07/08/2013 09:15:09
Registering Administrator: $$CONFIG_MANAGER$$
Bytes Received Last Session: 362
Bytes Sent Last Session: 507
Duration of Last Session: 0.00
Pct. Idle Wait Last Session: 0.00
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Grace Deletion Period: 5
Managing profile:
Server Password Set: Yes
Server Password Set Date/Time: 07/08/2013 09:15:09
Days Since Server Password Set: 1
Invalid Sign-on Count for Server: 0
Virtual Volume Password Set: No
Virtual Volume Password Set Date/Time: (?)
Days Since Virtual Volume Password Set: (?)
Invalid Sign-on Count for Virtual Volume Node: 0
Validate Protocol: No
Version: 7
Release: 1
Level: 0.0
Role(s): Replication
SSL: No
Session Security: Strict
Transport Method: TLS 1.2
```

Field descriptions

Server Name

The name of the server.

Comm. Method

The communication method that is used to connect to the server.

Transfer Method

The method that is used for server-to-server data transfer.

High-level Address

The IP address (in dotted decimal format) of the server.

Low-level Address

The port number of the server.

Description

The server description.

Allow Replacement

Specifies whether a server definition on a managed server can be replaced with a definition from a configuration manager.

Node Name

The name of the client node.

Last Access Date/Time

The last date and time that the client node accessed the server.

Days Since Last Access

The number of days since the client node accessed the server.

Compression

The type of compression that is completed by IBM Spectrum Protect™ on client files.

Archive Delete Allowed?

Specifies whether the client node can delete its own archive files. A value of (?) denotes that this field is not set and does not apply to this definition.

URL

The URL used to access this server from a web browser-based interface.

Registration Date/Time

The date and time that the client node was registered.

Registering Administrator

The name of the administrator that registered the client node.

Bytes Received Last Session

The number of bytes received by the server during the last client node session.

Bytes Sent Last Session

The number of bytes sent to the client node.

Duration of Last Session

The length of the last client node session, in seconds.

Pct. Idle Wait Last Session

The percentage of the total session time during which the client did not complete any functions.

Pct. Comm. Wait Last Session

The percentage of the total session time that the client waited for a response from the server.

Pct. Media Wait Last Session

The percentage of the total session time that the client waited for a removable volume to be mounted.

Grace Deletion Period

The number of days an object remains on the target server after it is marked for deletion.

Managing Profile

The profile from which the managed server got the definition of this server.

Server Password Set

Specifies whether the password for the server is set.

Server Password Set Date/Time

Specifies when the password for the server is set.

Days since Server Password Set

The number of days since the server password was set.

Invalid Sign-on count for Server

The maximum number of invalid sign-on attempts that the server can accept.

Virtual Volume Password Set

Specifies whether the password used to log on to the target server is set.

Virtual Volume Password Set Date/Time

Specifies when the password for virtual volume support is set.

Days Since Virtual Volume Password Set

The number of days since the password for virtual volume support was set.

Invalid Sign-on Count for Virtual Volume Node

The maximum number of invalid sign-on attempts that are accepted on the target server.

Validate Protocol (deprecated)

Specifies whether the storage agent has the data validation function enabled. This field is deprecated.

Version

The software version of the IBM Spectrum Protect server.

Release

The software release of the IBM Spectrum Protect server.

Level

The software level of the IBM Spectrum Protect server.

Role(s)

The role of the server. For example, one of the roles that the server is used for is replication.

SSL

Specifies whether Secure Sockets Layer (SSL) communication is used.

Session Security

Specifies the level of session security that is enforced for the server. Values can be STRICT or TRANSITIONAL.

Transport Method

Specifies the transport method that was last used for the specified server. Values can be TLS 1.2, TLS 1.1, or NONE. A question mark (?) is displayed until a successful authentication is completed.

Related commands

Table 1. Commands related to QUERY SERVER

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DEVCLASS	Deletes a device class.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
DELETE SERVER	Deletes the definition of a server.
AIX Linux Windows PROTECT STGPOOL	AIX Linux Windows Protects a directory-container storage pool.
QUERY NODE	Displays partial or complete information about one or more clients.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET REPLSERVER	Specifies a target replication server.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE NODE	Changes the attributes that are associated with a client node.
UPDATE SERVER	Updates information about a server.

QUERY SERVERGROUP (Query a server group)

Use this command to display information about server groups and group members.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-QUERY SERVERGroup--+-----+----->>
                        .*-----
                        '-group_name-'
```

Parameters

group_name

Specifies the server group to query. This parameter is optional. You can use wildcard characters to specify this name.

Example: List server groups

From a managed server, query all server groups. Field descriptions for field descriptions.

```
query servergroup *
```

```
Server Group  Members      Description      Managing Profile
-----
```

ADMIN_GROUP	SERVER_A	Headquarters	ADMIN_INFO
	SERVER_B		
	SERVER_C		
	SERVER_D		

Field descriptions

Server Group

The name of the server group.

Members

The group members.

Description

The description of the server group.

Managing Profile

The profile or profiles to which the managed server subscribed to get the definition of the server groups.

Related commands

Table 1. Commands related to QUERY SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVER	Displays information about servers.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

QUERY SESSION (Query client sessions)

Use this command to display information about administrative, node, and server sessions.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query SEssion--+-----+----->
                    '-sessnum-'

>--+-----+----->
    '-MINTIMethreshold---minutes-'

>--+-----+----->
    '-MAXTHROUGHput---kilobytes_per_second-'

    .-Format---Standard----.  .-Type---*-----
>--+-----+----->
    '-Format---+Standard+-'  '-Type---+Admin--+'
                    '-Detailed-'          +-Node---+
                                        '-Server-'

    .-CLIENTName---*-----
>--+-----+----->>
    '-CLIENTName-----client_name---
```

Parameters

sessnum

Specifies the number of the administrative or client node session to query. This parameter is optional. If you do not specify a value for this parameter, all sessions display.

MINTIMethreshold

Specifies to display sessions that had at least this number of minutes elapse from the time the client sent data to the server for storage. This parameter is optional. The minimum number of minutes is 1. The maximum number of minutes is 99999999.

MAXTHRoughput

Specifies to display sessions that are transferring data at a rate less than this number of kilobytes per second. This parameter is optional. The minimum number of kilobytes per second is 0. The maximum number of kilobytes per second is 99999999.

Format

Specifies how the information displays. This parameter is optional. The default value is STANDARD. The following values are possible:

Standard

Specifies that partial information displays for the session.

Detailed

Specifies that complete information displays for the session.

Type

Specifies the type of sessions to include in the query results. If you do not specify a value for this parameter, all types of sessions are queried. This parameter is optional. You can specify one of the following values:

Admin

Specifies that administrative sessions are displayed.

Node

Specifies that node sessions are displayed.

Server

Specifies that server sessions are displayed.

CLIENTName

Specifies the name of an administrator, client node, or server to be queried. You can specify one or more names. You can also specify node groups and proxy nodes. If you specify multiple names, separate the names with commas; use no intervening spaces. You can use wildcard characters with node names but not with node group names. The parameter is optional.

During node replication, the client name on the target server is displayed as *node_name (server_name)*, where *node_name* is the node whose data is being replicated, and *server_name* is the name of the source server. You can specify either the node name or the server name in the CLIENTName parameter to display the replication sessions.

Example: List active client node sessions

Display information about all administrative and client node sessions that are communicating with the server. See Field descriptions for field descriptions.

```
query session
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
4	TCP/IP	Run	0 S	1.4 K	162	Admin	WinNT	ADMIN

Example: Display detailed information about active client node sessions

Display detailed information about all administrative and client node sessions that are communicating with the server. See Field descriptions for field descriptions.

```
query session format=detailed
```

```
Sess Number: 4
Comm. Method: Tcp/Ip
Sess State: Run
Wait Time: 0 S
Bytes Sent: 1.4 K
Bytes Recvd: 162
Sess Type: Admin
```

```
Platform: WinNT
Client Name: ADMIN
Media Access Status:
User Name:
Date/Time First Data Sent:
Proxy By Storage Agent:
Actions:
Failover Mode: No
```

Field descriptions

Sess Number

Specifies a unique session identification number that is assigned by the server.

Comm. Method

Specifies the method that is used by the client to communicate with the server.

Sess State

Specifies the current communications state of the server. The following states are possible:

End

The session is ending (session resources are released).

IdleW

Waiting for client's next request (session is idle).

MediaW

The session is waiting for access to a sequential access volume.

RecvW

Waiting to receive an expected message from the client.

Run

The server is running a client request (and not waiting to send data).

SendW

The server is waiting to send data to the client (waiting for data to be delivered to the client node that was already sent).

SSLiW

The session is waiting for Secure Sockets Layer (SSL) initialization to complete.

Start

The session is starting (authentication is in progress).

Wait Time

Specifies the amount of time (seconds, minutes, or hours) the server is in the current state shown.

Bytes Sent

Specifies the number of bytes of data that is sent to the client node since the session was initiated.

Bytes Recvd

Specifies the number of bytes of data that is received from the client node since the session was initiated.

Sess Type

Specifies the type of session in process: ADMIN for an administrative session, NODE for a client node session, or SERVER. SERVER specifies the server starts a session and initiates server-to-server operations such as central configuration, library sharing, and storage agent sessions.

Platform

Specifies the type of operating system that is associated with the client.

Client Name

Specifies the name of the client node or the administrator.

For node replication sessions, the client name is updated to *node_name (server_name)* on the target server after data transfer starts.

Media Access Status

Specifies the type of media wait state. When a session is in a media wait state, this field displays a list of all mount points and sequential volumes for the session. The list of mount points specifies the device class and the associated storage pool. The list of volumes specifies the primary storage pool volumes in addition to any copy storage pool and active-data pool volumes along with their assigned storage pool.

The server allows multiple sessions to read and one session to write to a volume concurrently in a storage pool that is associated with the FILE or CENTERA device type. As a result, a volume in a storage pool with a device type of FILE or CENTERA can appear as the current volume for more than one session.

Proxy by Storage Agent

Specifies the storage agent that is the proxy for LAN-free data movement for the node.

User Name

Specifies the user ID of the node, on a multi-user system, that connects to the server when it is not the same system user who originally connected to the server.

Date/Time First Data Sent

Specifies the date and time that the client first sent data to the server for storage.

Actions

Displays a list of actions that are performed during the session. An action is listed only once, even if the action occurs multiple times during a session. The following actions are possible:

BkIns

One or more backup objects were stored on the server. The operation might have been an incremental backup or a selective backup.

BkUpd

One or more attributes were updated for a backup object that is stored on the server.

BkDel

One or more backup objects that are stored on the server are deleted.

BkRebind

One or more backup objects that are stored on the server are bound to a different management class.

NoQueryRestore

A no-query restore operation was initiated from the client to restore backed-up files from the server to the client system.

ArIns

One or more archive objects were stored on the server.

ObjRtrv

One or more files were retrieved from the server. This might have been to retrieve archive files, or to restore backup data (except for backup data from a no-query restore operation).

MigIns

One or more files are migrated and stored on the server by IBM Spectrum Protect™ for Space Management (HSM client).

MigDel

One or more space-managed files that were stored on the server are deleted.

MigRebind

One or more space-managed files that are stored on the server are bound to a different management class.

MigRecall

One or more space-managed files that are stored on the server are recalled.

MigUpd

The attributes for one or more space-managed files that are stored on the server are updated.

FSAdd

The client node added one or more new file spaces to server storage.

FSUpd

The client node updated attributes for one or more file spaces that are defined to the server.

DefAuth

A SET ACCESS command is processed by the client node, which caused an authorization rule for access to the client node's data to be added.

Failover Mode

Specifies whether the client session was started in failover mode. The following values are possible:

Force

The FORCEFAILOVER flag is specified on the client and the session is forced into failover mode.

Yes

The client session was started in failover mode.

No

The client session was not started in failover mode.

Related commands

Table 1. Command related to QUERY SESSION

Command	Description
CANCEL SESSION	Cancels active sessions with the server.

QUERY SHREDSTATUS (Query shredding status)

Use this command to display information about data waiting to be shredded.

Privilege class

To issue this command you must have administrator privilege.

Syntax

```
>>-QUERY SHREDstatus-.Format---Standard-----><
                        +-----+-----+-----+-----+
                        '-Format---+Standard-+-'
                        '-Detailed-'
```

Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed. This is the default.

Detailed

Specifies that complete information is displayed.

Example: Display summary shredding information

Show partial information about data shredding on the server. See Field descriptions for field descriptions.

```
query shredstatus
```

Shredding	Objects
Active	Awaiting
	Shred
-----	-----
NO	4

Example: Display detailed shredding information

Display detailed information about data shredding on the server. See Field descriptions for field descriptions.

```
query shredstatus format=detailed
```

Shredding	Objects	Occupied	Data Left
Active	Awaiting	Space	To Shred
	Shred	(MB)	(MB)
-----	-----	-----	-----
NO	4	182	364

Field descriptions

Shredding Active

Indicates whether or not the server is actively shredding data at this time.

Objects Awaiting Shred

The number of objects currently waiting to be shredded.

Occupied Space (MB)

The amount of server storage space occupied by the objects currently waiting to be shredded, in megabytes. This is the amount of space that will become available when the objects are shredded.

Data Left to Shred (MB)

The amount of data that still needs to be shredded.

Related commands

Table 1. Commands related to QUERY SHREDSTATUS

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
EXPORT NODE	Copies client node information to external media or directly to another server.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
QUERY STGPOOL	Displays information about storage pools.
SETOPT	Updates a server option without stopping and restarting the server.
SHRED DATA	Manually starts the process of shredding deleted data.
UPDATE STGPOOL	Changes the attributes of a storage pool.

QUERY SPACETRIGGER (Query the space triggers)

Use this command to display the settings for storage pool space triggers.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query SPACETrigger--STG--+-----+----->
                        '-STGPOOL---storage_pool-'

.-Format----Standard----.
>--+-----+----->>
  '-Format---+Standard-+'
                        '-Detailed-'
```

Parameters

STG

Specifies a storage pool space trigger.

STGPOOL

Specifies one or more storage pools (using a wildcard) for which storage pool trigger information will be displayed. If STG is specified but STGPOOL is not, the default storage pool space trigger, if any, is displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display detailed settings for a storage pool space trigger

Issue this command:

```
query spacetrigger stg stgpool=archivepool format=detailed
```

AIX

```
STGPOOL Full Percentage: 50
STGPOOL Expansion Percentage: 20
STGPOOL Expansion prefix: /usr/tivoli/tsm/server/filevol/
STGPOOL: ARCHIVEPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/10/2004 11:59:59
```

Linux

```
STGPOOL Full Percentage: 50
STGPOOL Expansion Percentage: 20
STGPOOL Expansion prefix: /opt/tivoli/tsm/server/filevol/
STGPOOL: ARCHIVEPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/10/2004 11:59:59
```

Windows

```
STGPOOL Full Percentage: 50
STGPOOL Expansion Percentage: 20
STGPOOL Expansion prefix: c:\program files\tivoli\filevol\
STGPOOL: ARCHIVEPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/10/2004 11:59:59
```

Field descriptions

STGPOOL Full Percentage

The trigger utilization percentage at which IBM Spectrum Protect™ allocates more space for the storage pool.

STGPOOL Expansion Percentage

The percentage of space by which the storage pool should be expanded.

STGPOOL Expansion prefix

The prefix associated with the space trigger.

STGPOOL

The storage pool name associated with the query.

Last Update by (administrator)

The administrator who last updated the storage pool space trigger.

Last Update Date/Time

The date and time when the administrator last updated the storage pool space trigger.

Related commands

Table 1. Commands related to QUERY SPACETRIGGER

Command	Description
DEFINE SPACETRIGGER	Defines a space trigger to expand the space for a storage pool.
DELETE SPACETRIGGER	Deletes the storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

QUERY STATUS (Query system parameters)

Use the QUERY STATUS command to display information about system parameters.

Use this command for the following reasons:

- To display the service level of the server
- To display information about the general server parameters, such as those defined by the SET commands
- To request information about client sessions, such as the availability of the server, password authentication, accounting settings, or the retention period for the information that is retained in the activity log
- To display information about the central scheduler, such as the central scheduling mode of the server
- To display the maximum number of repeated attempts that are allowed after a failed attempt to run a scheduled command
- To display whether subfiles can be backed up to this server, as indicated by the SET SUBFILE command
- To display information about a target replication server

- To display licensing information

Tip: To display information about a target replication server, you must issue the command from the target replication server.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query STATUS-----<<
```

Parameters

None.

Example: Query the status of a configuration manager

Display general information about server parameters. The command is run from a configuration manager. For descriptions of displayed fields, see Field descriptions.

```
query status
```

AIX

```

Server Name: SETSHOT
Server host name or IP address: setshot
  Server TCP/IP port number: 1500
    Crossdefine: On
      Server Password Set: Yes
        Server Installation Date/Time: 2016-07-08, 09:45:53
          Server Restart Date/Time: 2016-10-10, 05:38:49
            Authentication: Off
              Password Expiration Period: 9,999 Day(s)
                Invalid Sign-on Attempt Limit: 0
                  Minimum Password Length: 8
                    Registration: Closed
                      Subfile Backup: Client
                        Availability: Enabled
                          Inbound Sessions Disabled:
                            Outbound Sessions Disabled:
                              Accounting: Off
                                Activity Log Retention: 30 Day(s)
                                  Activity Log Number of Records: 222919
                                    Activity Log Size: 6 M
                                      Activity Summary Retention Period: 30 Day(s)
                                        License Audit Period: 30 Day(s)
                                          Last License Audit: 2016-10-21, 07:40:20
                                            Server License Compliance: Valid
                                              Central Scheduler: Active
                                                Maximum Sessions: 300
                                                  Maximum Scheduled Sessions: 75
                                                    Event Record Retention Period: 14 Day(s)
                                                      Client Action Duration: 5 Day(s)
                                                        Schedule Randomization Percentage: 25
                                                          Query Schedule Period: Client
                                                            Maximum Command Retries: Client
                                                              Retry Period: Client
                                                                Client-side Deduplication Verification Level: 0 %
                                                                  Scheduling Modes: Any
                                                                    Active Receivers: CONSOLE ACTLOG
                                                                      Configuration manager?: Off
                                                                        Refresh interval: 60
                                                                          Last refresh date/time:
                                                                            Context Messaging: On
                                                                              Table of Contents (TOC) Load Retention: 120 Minute(s)
                                                                                Machine Globally Unique ID: d4.cg.f6.ae.04.6e.11.e3.80.1f.00.21.5e.18.df.01
                                                                                  Archive Retention: Off
                                                                                    Database Directories: /TSMserver/DB1,/TSMserver/DB2

```

Total Space of File System (MB): 222,720.00
 Used Space on File System (MB): 47,780.74
 Free Space Available (MB): 174,939.26
 Encryption Strength: AES
 Client CPU Information Refresh Interval: 180
 Outbound Replication: Enabled
 Target Replication Server: POWER
 Default Replication Rule for Archive: ALL_DATA
 Default Replication Rule for Backup: ALL_DATA
 Default Replication Rule for Space Management: ALL_DATA
 Replication Record Retention Period: 30 Day(s)
 LDAP User:
 LDAP Password Set: No
 Default Authentication: Local
 Failover High Level Address:
 Scratchpad retention: 365 Day(s)
 Replication Recovery of Damaged Files: On
 SUR Occupancy (TB): 5.66
 SUR Occupancy Date/Time: 2016-10-10, 05:39:33
 Front-End Capacity (MB): 226,331
 Front-End Client Count: 6
 Front-End Capacity Date: 2016-10-13, 09:20:02
 Product Offering: IBM Spectrum Protect

Linux

Server Name: GOBI
 Server host name or IP address:
 Server TCP/IP port number: 1500
 Crossdefine: On
 Server Password Set: Yes
 Server Installation Date/Time: 2016-07-08, 11:29:03
 Server Restart Date/Time: 2016-11-10, 14:25:03
 Authentication: On
 Password Expiration Period: 90 Day(s)
 Invalid Sign-on Attempt Limit: 0
 Minimum Password Length: 8
 Registration: Closed
 Subfile Backup: No
 Availability: Enabled
 Inbound Sessions Disabled:
 Outbound Sessions Disabled:
 Accounting: Off
 Activity Log Retention: 30 Day(s)
 Activity Log Number of Records: 21346
 Activity Log Size: <1 M
 Activity Summary Retention Period: 30 Day(s)
 License Audit Period: 30 Day(s)
 Last License Audit: 2016-10-21, 23:27:23
 Server License Compliance: Valid
 Central Scheduler: Active
 Maximum Sessions: 500
 Maximum Scheduled Sessions: 250
 Event Record Retention Period: 14 Day(s)
 Client Action Duration: 5 Day(s)
 Schedule Randomization Percentage: 25
 Query Schedule Period: Client
 Maximum Command Retries: Client
 Retry Period: Client
 Client-side Deduplication Verification Level: 0 %
 Scheduling Modes: Any
 Active Receivers: CONSOLE ACTLOG
 Configuration manager?: Off
 Refresh interval: 60
 Last refresh date/time:
 Context Messaging: Off
 Table of Contents (TOC) Load Retention: 120 Minute(s)
 Machine Globally Unique ID: fc.e7.be.58.4a.a7.11.e0.8a.c8.e4.1f.13.34.11.e0
 Archive Retention Protection: Off
 Database Directories:
 /TSMdbspace1/gpcinst1,/TSMdbspace2/gpcinst1,/TSMdbspace3/gpcinst1
 Total Space of File System (MB): 302,379.84
 Used Space on File System (MB): 106,793.65
 Free Space Available (MB): 195,586.20
 Encryption Strength: AES

Client CPU Information Refresh Interval: 180
 Outbound Replication: Enabled
 Target Replication Server:
 Default Replication Rule for Archive: ALL_DATA
 Default Replication Rule for Backup: ALL_DATA
 Default Replication Rule for Space Management: ALL_DATA
 Replication Record Retention Period: 30 Day(s)
 LDAP User:
 LDAP Password Set: No
 Default Authentication: Local
 Failover High Level Address:
 Scratchpad retention: 365 Day(s)
 Replication Recovery of Damaged Files: Off
 SUR Occupancy (TB): 0.00
 SUR Occupancy Date/Time: 2016-10-10, 14:25:35
 Front-End Capacity (MB): 226,331
 Front-End Client Count: 6
 Front-End Capacity Date: 2016-10-13, 09:20:02
 Product Offering: IBM Spectrum Protect

Windows

Server Name: EXCELSIOR
 Server host name or IP address: excelsior.storage.
 newyork.example.com
 Server TCP/IP port number: 1500
 Crossdefine: On
 Server Password Set: Yes
 Server Installation Date/Time: 2016-07-08, 18:02:50
 Server Restart Date/Time: 2016-11-10, 11:48:32
 Authentication: On
 Password Expiration Period: 90 Day(s)
 Invalid Sign-on Attempt Limit: 0
 Minimum Password Length: 8
 Registration: Closed
 Subfile Backup: No
 Availability: Enabled
 Inbound Sessions Disabled:
 Outbound Sessions Disabled:
 Accounting: On
 Activity Log Retention: 30 Day(s)
 Activity Log Number of Records: 1346376
 Activity Log Size: 37 M
 Activity Summary Retention Period: 30 Day(s)
 License Audit Period: 30 Day(s)
 Last License Audit: 2016-10-21, 17:05:16
 Server License Compliance: Valid
 Central Scheduler: Active
 Maximum Sessions: 25
 Maximum Scheduled Sessions: 12
 Event Record Retention Period: 14 Day(s)
 Client Action Duration: 5 Day(s)
 Schedule Randomization Percentage: 25
 Query Schedule Period: Client
 Maximum Command Retries: Client
 Retry Period: Client
 Client-side Deduplication Verification Level: 0 %
 Scheduling Modes: Any
 Active Receivers: CONSOLE ACTLOG
 NTEVENTLOG
 Configuration manager?: Off
 Refresh interval: 60
 Last refresh date/time:
 Context Messaging: Off
 Table of Contents (TOC) Load Retention: 120 Minute(s)
 Machine Globally Unique ID: e9.3e.f1.70.ff.c5.11.e2.
 a5.67.5c.f3.fc.0c.5e.60
 Archive Retention Protection: Off
 Database Directories: e:\Server1\TSMDBdir
 Total Space of File System (MB): 102,270.00
 Used Space on File System (MB): 22,032.79
 Free Space Available (MB): 80,237.20
 Encryption Strength: AES
 Client CPU Information Refresh Interval: 180

```

        Outbound Replication: Enabled
        Target Replication Server: EXPLORER
    Default Replication Rule for Archive: ALL_DATA
    Default Replication Rule for Backup: ALL_DATA
Default Replication Rule for Space Management: ALL_DATA
    Replication Record Retention Period: 30 Day(s)
        LDAP User: cn=excelsior_ldapadmin,ou=excelsior,
                   ou=John Doe,dc=tsmadldap,dc=storage,
                   dc=newyork, dc=example,dc=com

        LDAP Password Set: Yes
        Default Authentication: LDAP
    Failover High Level Address:
        Scratchpad retention: 365 Day(s)
    Replication Recovery of Damaged Files: On
        SUR Occupancy (TB): 8.98
        SUR Occupancy Date/Time: 2016-10-10, 11:49:27
    Front-End Capacity (MB): 226,331
    Front-End Client Count: 6

```

Windows

```

    Front-End Capacity Date: 2016-10-13, 09:20:02
    Product Offering: IBM Spectrum Protect

```

Field descriptions

- Server Name**
Specifies the name of the server.
- Server host name or IP address**
Specifies the server TCP/IP address.
- Server TCP/IP port number**
Specifies the server port address.
- Crossdefine**
Specifies whether another server that is running the DEFINE SERVER command automatically defines itself to this server.
See the SET CROSSDEFINE command.
- Server Password Set**
Specifies whether the password was set for the server.
- Server Installation Date/Time**
Specifies the date and time when the server was installed.
- Server Restart Date/Time**
Specifies the last date and time when the server was started.
- Authentication**
Specifies whether password authentication is set on or off.
- Password Expiration Period**
Specifies the period, in days, after which the administrator or client node password expires.
- Invalid Sign-on Attempt Limit**
Specifies the number of invalid sign-on attempts before a node is locked.
- Minimum Password Length**
Specifies the minimum number of characters for the password. This value does not apply to configurations where an LDAP server is used.
- Registration**
Specifies whether client node registration is open or closed.
- Subfile Backup**
Specifies whether subfiles can be backed up to this server, as indicated by the SET SUBFILE command.
- Availability**
Specifies whether the server is enabled or disabled.
- Inbound Sessions Disabled**
Specifies the names of servers from which server-to-server communications are not allowed. To enable inbound server sessions, use the ENABLE SESSIONS command.
- Outbound Sessions Disabled**
Specifies the names of servers to which server-to-server communications are not allowed. To enable outbound server sessions, use the ENABLE SESSIONS command.
- Accounting**
Specifies whether an accounting record is generated at the end of each client node session.
- Activity Log Retention**
Specifies the number of days information is retained in the activity log, or the size of the log.

Activity Log Number of Records
Specifies the number of records in the activity log.

Activity Log Size
Specifies the size of the activity log.

Activity Summary Retention Period
Specifies the number of days information is retained in the SQL activity summary table.

License Audit Period
Specifies the period, in days, after which the license manager automatically audits the IBM Spectrum Protect™ license. Additional licensing information is available by using the QUERY LICENSE command.

Last License Audit
Specifies the date and time when the last license audit occurred. Additional licensing information is available by using the QUERY LICENSE command.

Server License Compliance
Specifies whether the server is in compliance (Valid) or out of compliance (Failed) with the license terms. Use the QUERY LICENSE command to see what factors are causing the server to fail to comply with the license terms.

Central Scheduler
Specifies whether central scheduling is running (active or inactive).

Maximum Sessions
Specifies the maximum number of client/server sessions.

Maximum Scheduled Sessions
Specifies the maximum number of client/server sessions available for processing scheduled work.

Event Record Retention Period
Specifies the number of days central scheduler event records are retained.

Client Action Duration
Specifies the duration of the period during which the client processes the schedule that is defined with the DEFINE CLIENTACTION command.

Schedule Randomization Percentage
Specifies how much of the startup window is used for running scheduled events in client-polling mode.

Query Schedule Period
Specifies the frequency with which clients poll the server to obtain scheduled work, in client-polling mode. If the value in this field is Client, the polling frequency is determined by the client node.

Maximum Command Retries
Specifies the maximum number of times that a client scheduler tries to run a scheduled command after a failed attempt. If the value in this field is Client, the client node determines the maximum number.

Retry Period
Specifies the number of minutes between failed attempts by the client scheduler to contact the server or to run a scheduled command. If the value in this field is Client, the client node determines the number of minutes.

Client-side Deduplication Verification Level
Specifies a percentage of extents to be verified by the IBM Spectrum Protect server. The extents are created during client-side data deduplication.

Scheduling Modes
Specifies the central scheduling modes that are supported by the server.

Active Receivers
Specifies the receivers for which event logging began.

Configuration manager?
Specifies whether the server is a configuration manager.

Refresh interval
Specifies the interval that elapses before the managed server requests a refresh of any changes from a configuration manager.

Last refresh date/time
If the server is a managed server, specifies the date and time of the last successful refresh of configuration information from the configuration manager.

Context Messaging
Specifies whether context messaging is enabled or disabled.

Table of Contents (TOC) Load Retention
Specifies the approximate number of minutes that unreferenced TOC data is retained in the database.

Machine Globally Unique ID
The globally unique identifier (GUID) as of the last time that the server was started. This GUID identifies the host system to which the current server belongs.

Archive Retention Protection
Specifies whether archive data retention protection is activated or deactivated.

Database Directories

Specifies the locations of the database directories.

Total Space of File System (MB)
Specifies the total size of the file system.

Used Space on File System (MB)
Specifies the amount of space that is in use on the file system.

Free Space Available (MB)
Specifies the amount of space that is available.

Encryption Strength
Indicates data encryption strength: AES or DES.

Client CPU Information Refresh Interval
Specifies the number of days that elapse between client scans for CPU information that is used for PVU estimation.

Outbound Replication
Specifies whether replication processing is enabled or disabled. If outbound replication is disabled, new replication processes cannot start on the server.

Target Replication Server
Specifies the name of the server that is the target for node replication operations. If a target replication server does not exist, this field is blank.

Default Replication Rule for Archive
Specifies the server replication rule that applies to archive data. The following values are possible:

ALL_DATA
Replicates archive data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY
Replicates archive data. The data is replicated with a high priority.

NONE
Archive data is not replicated.

Default Replication Rule for Backup
Specifies the server replication rule that applies to backup data. The following values are possible:

ALL_DATA
Replicates active and inactive backup data. The data is replicated with a normal priority.

ACTIVE_DATA
Replicates only active backup data. The data is replicated with a normal priority.
Attention: If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the FORCERECONCILE=YES parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY
Replicates active and inactive backup data. The data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY
This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority.

NONE
Backup data is not replicated.

Default Replication Rule for Space Management
Specifies the server replication rule that applies to space-managed data. The following values are possible:

ALL_DATA
Replicates space-managed data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY
Replicates space-managed data. The data is replicated with a high priority.

NONE
Space-managed data is not replicated.

Replication Record Retention Period

Specifies the number of days that replication history records are retained in the database of the source replication server.

LDAP User

Specifies the user ID that is named in the SET LDAPUSER command. This user ID can issue administrative commands on the namespace that is reserved for IBM Spectrum Protect on the LDAP directory server.

LDAP Password Set

This output field shows if a password is defined for the user ID that is named in the SET LDAPUSER command. The values are YES and NO. If YES, the user ID that is named in the SET LDAPUSER command can issue administrative commands on the LDAP namespace that is reserved for IBM Spectrum Protect. If NO, issue the SET LDAPPASSWORD command to set the password for the user ID that is named in the SET LDAPUSER command.

Default Authentication

Specifies the default password authentication method: LOCAL or LDAP.

Authentication Target	Authentication Method
IBM Spectrum Protect server	LOCAL
LDAP directory server	LDAP

When you issue the SET DEFAULTAUTHENTICATION command, you define the resulting authentication method for all REGISTER ADMIN and REGISTER NODE commands. The default is LOCAL.

Failover High Level Address

Specifies the high-level address for the failover server that is used by the client. Client restore operations fail over to this high-level address when the interface that is used by the client is different from the interface that is used by replication.

Scratchpad retention

Specifies the number of days for which scratch pad entries are retained since they were last updated.

Replication Recovery of Damaged Files

Specifies whether node replication is enabled to recover damaged files from a target replication server. This is a system-side setting. If ON is specified, the node replication process can be configured to detect damaged files on a source replication server and replace them with undamaged files from a target replication server. If OFF is specified, damaged files are not recovered from a target replication server.

SUR Occupancy (TB)

If you have an IBM Spectrum Protect Suite (SUR) license, this field specifies the SUR occupancy on the server. The *SUR occupancy* is the amount of space that is used to store data that is managed by the IBM Spectrum Protect products that are included in the SUR bundle.

SUR Occupancy Date/Time

Specifies the date and time when SUR occupancy data was last collected.

Front-End Capacity (MB)

Specifies the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems. This value is used for the front-end licensing model.

Front-End Client Count

Specifies the number of clients that reported capacity usage based on the front-end licensing model.

Front-End Capacity Date

Specifies the date and time when front-end capacity data was last collected.

Product Offering

Specifies a product offering.

Value specified by the SET PRODUCTOFFERING command	Value shown in the QUERY STATUS command output
ENTry	IBM Spectrum Protect Entry
DATARet	IBM Spectrum Protect for Data Retention
BASIC	IBM Spectrum Protect
EE	IBM Spectrum Protect Extended Edition
SUIte	IBM Spectrum Protect Suite
SUITECloud	IBM Spectrum Protect Suite - IBM Cloud Object Storage Option
SUITEEntry	IBM Spectrum Protect Suite Entry
SUITEArchive	IBM Spectrum Protect Suite - Archive
SUITEProtectier	IBM Spectrum Protect Suite - ProtecTier
SUITEFrontend	IBM Spectrum Protect Suite - FrontEnd

Value specified by the SET PRODUCTOFFERING command	Value shown in the QUERY STATUS command output
SUITEENTRYFrontend	IBM Spectrum Protect Suite Entry - FrontEnd
CLEAR	NULL

Related commands

Table 1. Commands related to QUERY STATUS

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE REPLICATION	Prevents outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing IBM Spectrum Protect but permits existing sessions to continue.
ENABLE REPLICATION	Allows outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY LICENSE	Displays information about licenses and audits.
SET ACCOUNTING	Specifies whether accounting records are created at the end of each client session.
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.
SET CONTEXTMESSAGING	Specifies to turn on context messaging to debug an ANR9999D message.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET CROSSDEFINE	Specifies whether to cross define servers.
SET DEDUPVERIFICATIONLEVEL	Specifies the percentage of extents verified by the server during client-side deduplication.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET EVENTRETENTION	Specifies the number of days to retain records for scheduled operations.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.
SET MAXSCHEDESESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
SET PRODUCTOFFERING	Set the product offering licensed to your enterprise.
SET QUERYSCHEDPERIOD	Specifies the frequency for clients to obtain scheduled work, in client-polling mode.
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.

Command	Description
SET REPLRECOVERDAMAGED	Specifies whether node replication is enabled to recover damaged files from a target replication server.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.
SET SCHEDMODES	Specifies the central scheduling mode for the server.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERPASSWORD	Specifies the server password.
SET SUMMARYRETENTION	Specifies the number of days to retain information for the activity summary table.
SET TOCLOADRETENTION	Specifies the number of minutes to retain information for unreferenced TOC sets.

QUERY STATUSTHRESHOLD (Query status monitoring thresholds)

Use this command to display information about status monitoring thresholds.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-*-----
>>-Query STAtusthreshold----->
      '-threshold_name-'

      .-Format----Standard----.
>--+-----+----->
      '-Format----+Standard-+-' '-Activity----activity-'
          '-Detailed-'

>--+-----+----->
      '-Condition----+EXists-+-' '-Value----value_name-'
          +-GT-----+
          +-GE-----+
          +-LT-----+
          +-LE-----+
          '-Equal--'

>--+-----+----->>
      '-Status----+Normal-+-'
          +-Warning-+
          '-Error--'

```

Parameters

threshold_name

Specifies the threshold name. The name cannot exceed 48 characters in length.

Format

Specifies how the information is displayed. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified status thresholds.

Detailed

Specifies that complete information is displayed for the specified status thresholds.

activity

Specifies the activity for which you want to display status indicators. If you do not specify a value, information is displayed for all activities. For a list of activities, see the DEFINE STATUSTHRESHOLD command.

Condition

Restricts the output to only those matching the specified value. Possible values are:

EXists

Displays status thresholds where the condition equals EXISTS.

GT

Displays status thresholds where the condition equals GT.

GE

Displays status thresholds where the condition equals GE.

LT

Displays status thresholds where the condition equals LT.

LE

Displays status thresholds where the condition equals LE.

EQual

Displays status thresholds where the condition equals EQUAL.

Value

Displays thresholds that have the specified value. If you do not specify a value, information is displayed for all values. You can specify an integer from 0 to 9223372036854775807.

Status

Displays status thresholds that have the specified status value. If you do not specify a value, information is displayed for all values. Possible values are:

Normal

Displays the status thresholds that have a normal status value.

Warning

Displays the status thresholds that have a warning status value.

Error

Displays the status thresholds that have an error status value.

QUERY status threshold

Query all status thresholds by issuing the following command:

```
query statusthreshold
```

Threshold Name	Activity Name	Condition Name	Value	Report Status
ACTIVELOGCHECK	ACTIVE LOG UTILIZATION (%)	>	90	ERROR
AVGSTGPLW	AVERAGE STORAGE POOL UTILIZATION (%)	>	85	WARNING
AVGSTGPLE	AVERAGE STORAGE POOL UTILIZATION (%)	>	90	ERROR

Query status thresholds and display detailed format

Query status thresholds and display the output in detailed format, by issuing the following command:

```
query statusthreshold f=d

Threshold Name: ACTIVELOGCHECK
Activity Name: ACTIVE LOG UTILIZATION (%)
Condition Name: >
Value: 90
Report Status: ERROR
Server Name: TSMAWP24

Threshold Name: AVGSTGPLW
Activity Name: AVERAGE STORAGE POOL UTILIZATION (%)
Condition Name: >
Value: 85
Report Status: WARNING
Server Name: TSMAWP24

Threshold Name: AVGSTGPLE
Activity Name: AVERAGE STORAGE POOL UTILIZATION (%)
Condition Name: >
Value: 95
Report Status: ERROR
Server Name: TSMAWP24
```

Related commands

Table 1. Commands related to QUERY STATUSTHRESHOLD

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

QUERY STGPOOL (Query storage pools)

Use this command to display information about one or more storage pools. You can also use this command to monitor migration processes for storage pools.

Privilege class

Any administrator can issue this command.

Syntax

```

.*----- .-Format---Standard----.
>>-Query STGpool----->
'-pool_name-' '-Format---Standard--'
'-Detailed-'

.-Pooltype---ANY-----
>-----<
'-Pooltype---ANY-----'
+-Primary-----+
+-Copy-----+
+-COPYCONtainer-+
'-ACTIVEdata----'

```

Parameters

pool_name

Specifies the storage pool to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all storage pools are displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Specify one of the following values:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Pooltype

Specifies the type of storage pool to query. This parameter is optional. The default value is ANY. Specify one of the following values:

ANY

Query primary storage pools, copy storage pools, and active-data pools.

Primary

Query only primary storage pools.

COPY

Query only copy storage pools.

COPYCONtainer

Query only container-copy storage pools.

ACTIVEdata

Query only active-data storage pools.

Example: Display detailed random-access disk storage pool information

Tip: In the examples of detailed output, some fields are blank because the item does not apply in the specified environment. Display details for a storage pool that is named DISKPOOL. See Field descriptions for field descriptions.

```

query stgpool diskpool format=detailed

Storage Pool Name: DISKPOOL
Storage Pool Type: Primary
Device Class Name: DISK
Storage Type: DEVCLASS
Cloud Type:
Cloud URL:
Cloud Identity:
Cloud Location:
Estimated Capacity: 66 G
Space Trigger Util: 0.0
Pct Util: 0.0
Pct Migr: 3.1
Pct Logical: 100.0
High Mig Pct: 90
Low Mig Pct: 70
Migration Delay: 0
Migration Continue: Yes

```

```

Migration Processes: 1
Reclamation Processes: 1
  Next Storage Pool:
  Reclaim Storage Pool:
Maximum Size Threshold: No Limit
  Access: Read/Write
  Description:
  Overflow Location:
Cache Migrated Files?:
  Collocate?: Group
  Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 32
Number of Scratch Volumes Used: 1
Delay Period for Container Reuse: 1 Day(s)
  Migration in Progress?: No
  Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
  Reclamation in Progress?: No

Last Update by (administrator): SERVER_CONSOLE
  Last Update Date/Time: 01/03/2014 13:57:16
Storage Pool Data Format: Native
  Copy Storage Pool(s):
  Active Data Pool(s):
  Continue Copy on Error?: No
  CRC Data: Yes
  Reclamation Type: Threshold
  Overwrite Data when Deleted: 2 Time(s)
  Deduplicate Data?: No
Processes For Identifying Duplicates:
  Compressed:
  Deduplication Savings:
  Compression Savings:
  Total Space Saved:
  Auto-copy Mode: Client
Contains Data Deduplicated by Client?: No
  Maximum Simultaneous Writers:
  Protect Processes:
  Protection Storage Pool:
  Protect Local Storage Pool(s):
  Reclamation Volume Limit:

Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
  Deduplicate Requires Backup?:
  Encrypted:
  Pct Encrypted:
Cloud Space Allocated (MB):
Cloud Space Utilized (MB):
  Bucket Name:
  Local Estimated Capacity:
  Local Pct Util:
  Local Pct Logical:

```

Example: Display detailed sequential-access disk storage pool information

Display details for a storage pool that is named FILEPOOL. See Field descriptions for field descriptions.

```
query stgpool filepool format=detailed
```

```

Storage Pool Name: FILEPOOL
Storage Pool Type: Primary
Device Class Name: FILEC
  Storage Type: DEVCLASS
  Cloud Type:
  Cloud URL:
  Cloud Identity:
  Cloud Location:
Estimated Capacity: 66 G
Space Trigger Util: 0.0
  Pct Util: 0.0
  Pct Migr: 3.1
  Pct Logical: 100.0

```

```

        High Mig Pct: 90
        Low Mig Pct: 70
        Migration Delay: 0
        Migration Continue: Yes
        Migration Processes: 1
        Reclamation Processes: 1
        Next Storage Pool:
        Reclaim Storage Pool:
        Maximum Size Threshold: No Limit
        Access: Read/Write
        Description:
        Overflow Location:
        Cache Migrated Files?:
        Collocate?: Group
        Reclamation Threshold: 60
        Offsite Reclamation Limit:
        Maximum Scratch Volumes Allowed: 32
        Number of Scratch Volumes Used: 1
        Delay Period for Container Reuse: 1 Day(s)
        Migration in Progress?: No
        Amount Migrated (MB): 0.00

Elapsed Migration Time (seconds): 0
        Reclamation in Progress?: No
        Last Update by (administrator): SERVER_CONSOLE
        Last Update Date/Time: 01/02/2014 13:57:16
        Storage Pool Data Format: Native
        Copy Storage Pool(s):
        Active Data Pool(s):
        Continue Copy on Error?: No
        CRC Data: Yes
        Reclamation Type: Threshold
        Overwrite Data when Deleted:
        Deduplicate Data?: Yes
        Processes For Identifying Duplicates: 1
        Compressed:
        Deduplication Savings: 65,396 K (49.99%)
        Compression Savings:
        Total Space Saved: 65,396 K (49.99%)
        Auto-copy Mode: Client
        Contains Data deduplicated by Client?: Yes
        Maximum Simultaneous Writers:
        Protect Processes:
        Protection Storage Pool:
        Protect Local Storage Pool(s):
        Reclamation Volume Limit:
        Date of Last Protection to Remote Pool:
        Date of Last Protection to Local Pool:
        Deduplicate Requires Backup?:
        Encrypted:
        Pct Encrypted:
        Cloud Space Allocated (MB):
        Cloud Space Utilized (MB):
        Bucket Name:
        Local Estimated Capacity:
        Local Pct Util:
        Local Pct Logical:

```

Example: Display detailed sequential storage pool information

Display details for an active-data sequential storage pool that is named FILEPOOL that uses a FILE type device class. See Field descriptions for field descriptions.

```
query stgpool filepool format=detailed
```

```

Storage Pool Name: FILEPOOL
Storage Pool Type: Active-data
Device Class Name: FILEC
Storage Type: DEVCLASS
Cloud Type:
Cloud URL:
Cloud Identity:
Cloud Location:

```

```

Estimated Capacity: 0.0 M
Space Trigger Util: 0.0
  Pct Util: 0.0
  Pct Migr: 0.0
  Pct Logical: 0.0
  High Mig Pct: 90
  Low Mig Pct: 70
  Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
  Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold: No Limit
  Access: Read/Write
  Description:
  Overflow Location:
Cache Migrated Files?:
  Collocate?: Group
Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 99
Number of Scratch Volumes Used: 0
Delay Period for Container Reuse: 1 Day(s)
Migration in Progress?: No
  Amount Migrated (MB): 0.00

Elapsed Migration Time (seconds): 0
  Reclamation in Progress?: No
  Last Update by (administrator): SERVER_CONSOLE
  Last Update Date/Time: 01/02/2014 11:37:57
Storage Pool Data Format: Native
  Copy Storage Pool(s):
  Active Data Pool(s):
Continue Copy on Error?:
  CRC Data: Yes
  Reclamation Type: Threshold
Overwrite Data when Deleted:
Deduplicate Data?: Yes
Processes For Identifying Duplicates: 1
  Compressed:
  Deduplication Savings: 65,396 K (49.99%)
  Compression Savings:
  Total Space Saved: 65,396 K (49.99%)
  Auto-copy Mode:
Contains Data Deduplicated by Client?: No
  Maximum Simultaneous Writers:
  Protect Processes:
  Protection Storage Pool:
Protect Local Storage Pool(s):
  Reclamation Volume Limit:
Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
  Deduplicate Requires Backup?:
  Encrypted:
  Pct Encrypted:
Cloud Space Allocated (MB):
Cloud Space Utilized (MB):
  Bucket Name:
Local Estimated Capacity:
  Local Pct Util:
  Local Pct Logical:

```

Example: Display summary information for a specific storage pool

Display information for a storage pool that is named POOL1. See Field descriptions for field descriptions.

```
query stgpool pool1
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
POOL1	DISK	58.5 M	0.8	0.7	90	70	POOL2

Example: Display detailed 8 mm tape storage pool information

Display details for the storage pool named 8MMPOOL. See Field descriptions for field descriptions.

```
query stgpool 8mmpool format=detailed
```

```
Storage Pool Name: 8MMPOOL
Storage Pool Type: Primary
Device Class Name: 8MMTAPE
Storage Type: DEVCLASS
Cloud Type:
Cloud URL:
Cloud Identity:
Cloud Location:
Estimated Capacity: 0.0 M
Space Trigger Util: 0.0
Pct Util: 0.0
Pct Migr:
Pct Logical: 0.0
High Mig Pct: 90
Low Mig Pct: 70
Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold: 5 M
Access: Read/Write
Description: Main storage pool
Overflow Location: Room1234/Bldg31
Cache Migrated Files?:
Collocate?: No
Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 5
Number of Scratch Volumes Used: 3
Delay Period for Container Reuse: 1 Day(s)
Migration in Progress?: No
Amount Migrated (MB): 0.00

Elapsed Migration Time (seconds): 0
Reclamation in Progress?: No
Last Update by (administrator): ADMIN
Last Update Date/Time: 01/08/2014 06:55:45
Storage Pool Data Format: Native
Copy Storage Pool(s): COPYPOOL1
Active Data Pool(s): ACTIVEPOOL1 ACTIVEPOOL2
Continue Copy on Error?: Yes
CRC Data: Yes
Reclamation Type: Threshold
Overwrite Data when Deleted:
Deduplicate Data?: No
Processes For Identifying Duplicates:
Compressed:
Deduplication Savings:
Compression Savings:
Total Space Saved:
Compressed: No
Deduplication Savings:
Compression Savings:
Total Space Saved:
Auto-copy Mode: Client
Contains Data Deduplicated by Client?: No
Maximum Simultaneous Writers:
Protect Processes:
Protection Storage Pool:
Protect Local Storage Pool(s):
Reclamation Volume Limit:

Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
Deduplicate Requires Backup?:
Encrypted:
```



```
Pct Encrypted:
Cloud Space Allocated (MB):
Cloud Space Utilized (MB):
  Bucket Name:
Local Estimated Capacity:
  Local Pct Util:
  Local Pct Logical:
```

Example: Display detailed NAS2CLASS storage pool information

Display details for a storage pool, NAS2LIBPOOL. When you set up this storage pool, you set the data format to NETAPPDUMP. See Field descriptions for field descriptions.

```
query stgpool nas2libpool format=detailed
```

```
Storage Pool Name: NAS2
Storage Pool Name: NAS2LIBPOOL
Storage Pool Type: Primary
Device Class Name: NAS2CLASS
  Storage Type: DEVCLASS
  Cloud Type:
  Cloud URL:
  Cloud Identity:
  Cloud Location:
Estimated Capacity: 0.0 M
Space Trigger Util:
  Pct Util: 0.0
  Pct Migr:
  Pct Logical: 0.0
  High Mig Pct:
  Low Mig Pct:
  Migration Delay:
  Migration Continue:
  Migration Processes:
  Reclamation Processes:
  Next Storage Pool:
  Reclaim Storage Pool:
Maximum Size Threshold:
  Access: Read/Write
  Description:
  Overflow Location:
  Cache Migrated Files?:
  Collocate?: Group
  Reclamation Threshold:
  Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 50
  Number of Scratch Volumes Used: 0
Delay Period for Container Reuse: 1 Day(s)
  Migration in Progress?:
  Amount Migrated (MB):

Elapsed Migration Time (seconds):
  Reclamation in Progress?:
  Last Update by (administrator): SERVER_CONSOLE
  Last Update Date/Time: 01/02/2014 16:24:43
Storage Pool Data Format: NetApp Dump
  Copy Storage Pool(s):
  Active Data Pool(s):
  Continue Copy on Error?: No
  CRC Data: No
  Reclamation Type:
  Overwrite Data when Deleted:
  Deduplicate Data?: No
Processes For Identifying Duplicates:
  Compressed:
  Deduplication Savings:
  Compression Savings:
  Total Space Saved:
  Auto-copy Mode: Client
Contains Data Deduplicated by Client?: No
  Maximum Simultaneous Writers:
  Protect Processes:
  Protection Storage Pool:
```

```

Protect Local Storage Pool(s):
  Reclamation Volume Limit:

Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
  Deduplicate Requires Backup?:
    Encrypted:
      Pct Encrypted:
Cloud Space Allocated (MB):
Cloud Space Utilized (MB):
  Bucket Name:
Local Estimated Capacity:
  Local Pct Util:
  Local Pct Logical:

```

Example: Display detailed information for a directory-container storage pool that is used for data deduplication

Display details for a directory-container storage pool, DPOOL1. See Field descriptions for field descriptions.

```
query stgpool dpool1 format=detailed
```

```

Storage Pool Name: DPOOL1
Storage Pool Type: Primary
Device Class Name:
  Storage Type: Directory
  Cloud Type:
  Cloud URL:
  Cloud Identity:
  Cloud Location:
Estimated Capacity: 798 G
Space Trigger Util:
  Pct Util: 3.4
  Pct Migr:
  Pct Logical: 100.0
  High Mig Pct:
  Low Mig Pct:
  Migration Delay:
  Migration Continue:
  Migration Processes:
  Reclamation Processes:
  Next Storage Pool:
  Reclaim Storage Pool:
Maximum Size Threshold: No Limit
  Access: Read/Write
  Description:
  Overflow Location:
  Cache Migrated Files?:
  Collocate?:
  Reclamation Threshold:
  Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed:
  Number of Scratch Volumes Used:
Delay Period for Container Reuse: 1 Day(s)
  Migration in Progress?:
  Amount Migrated (MB):

Elapsed Migration Time (seconds):
  Reclamation in Progress?:
  Last Update by (administrator): SERVER_CONSOLE
  Last Update Date/Time: 01/02/2014 16:24:43
Storage Pool Data Format: Native
  Copy Storage Pool(s):
  Active Data Pool(s):
  Continue Copy on Error?:
  CRC Data: No
  Reclamation Type:
  Overwrite Data when Deleted:
  Deduplicate Data?: Yes
Processes For Identifying Duplicates:
  Compressed: Yes
Space Used for Protected Data: 1,599 M

```

```

        Total Pending Space: 100 M
        Deduplication Savings: 1,331 M (67.56%)
        Compression Savings: 194,805 K (29.82%)
        Total Space Saved: 1,521 M (77.22%)
        Auto-copy Mode:
Contains Data Deduplicated by Client?:
    Maximum Simultaneous Writers: No Limit
        Protect Processes:
    Protection Storage Pool: DPOOL2
Protect Local Storage Pool(s):
    Reclamation Volume Limit:

Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
    Deduplicate Requires Backup?:
        Encrypted:
            Pct Encrypted: 34.56%
    Cloud Space Allocated (MB):
    Cloud Space Utilized (MB):
        Bucket Name:
    Local Estimated Capacity:
        Local Pct Util:
        Local Pct Logical:

```

Example: Display detailed information for a cloud-container storage pool that is used for data deduplication

Display details for a cloud container storage pool, CPOOL1. See Field descriptions for field descriptions.

```
query stgpool cpool1 format=detailed
```

```

        Storage Pool Name: CPOOL1
        Storage Pool Type: Primary
        Device Class Name:
            Storage Type: CLOUD
            Cloud Type: SWIFT
            Cloud URL: http://localhost.local
        Cloud Identity: Bailey
        Cloud Location: ONPREMISE
    Estimated Capacity:
    Space Trigger Util:
        Pct Util:
        Pct Migr:
            Pct Logical: 0.0
        High Mig Pct:
        Low Mig Pct:
    Migration Delay:
    Migration Continue:
    Migration Processes:
    Reclamation Processes:
        Next Storage Pool:
        Reclaim Storage Pool:
    Maximum Size Threshold: No Limit
        Access: Read/Write
    Description:
    Overflow Location:
    Cache Migrated Files?:
        Collocate?:
    Reclamation Threshold:
    Offsite Reclamation Limit:
    Maximum Scratch Volumes Allowed:
    Number of Scratch Volumes Used:
        Delay Period for Volume Reuse: 1
    Migration in Progress?:
        Amount Migrated (MB):

Elapsed Migration Time (seconds):
    Reclamation in Progress?:
    Last Update by (administrator): CODY
        Last Update Date/Time: 2015-05-28, 10:47:52
    Storage Pool Data Format: Native
    Copy Storage Pool(s):

```

```

Active Data Pool(s):
Continue Copy on Error?:
    CRC Data: No
    Reclamation Type:
Overwrite Data when Deleted:
Deduplicate Data?: Yes
Processes For Identifying Duplicates:
    Compressed: Yes
    Deduplication Savings: 9,241 K (89.76%)
    Compression Savings: 1,033 K (98.81%)
    Total Space Saved: 10,274 K (99.79%)
    Auto-copy Mode:
Contains Data Deduplicated by Client?:
    Maximum Simultaneous Writers: No Limit
    Protect Processes:
    Protection Storage Pool:
Protect Local Storage Pool(s):
    Reclamation Volume Limit:

Date of Last Protection to Remote Pool:
Date of Last Protection to Local Pool:
Deduplicate Requires Backup?:
    Encrypted: Yes
    Pct Encrypted: 34.56%
Cloud Space Allocated (MB): 4,231
Cloud Space Utilized (MB): 4,231
    Bucket Name:
    Local Estimated Capacity: 168 G
    Local Pct Util: 0.1
    Local Pct Logical: 100.0

```

Field descriptions

Storage Pool Name

The name of the storage pool.

Storage Pool Type

The type of storage pool.

Device Class Name

The name of the device class that is assigned to the storage pool.

Storage Type

The type of storage that is defined for the storage pool. The following storage types can be shown:

DEVCLASS

The storage pool specifies a device class, which determines the type of device where data is stored.

DIRECTORY

The storage pool creates logical containers for data in file system directories.

CLOUD

The storage pool creates logical containers for data in a cloud environment.

Cloud Type

For cloud storage pools, the type of cloud platform.

Cloud URL

For cloud storage pools, the URL for accessing the on-premises private cloud or off-premises public cloud.

Cloud Identity

For cloud storage pools, the user ID for accessing the on-premises private cloud or off-premises public cloud.

Cloud Location

For cloud storage pools, indicates whether the cloud is an on-premises private cloud or off-premises public cloud.

Estimated Capacity

The estimated capacity of the storage pool in megabytes (M) or gigabytes (G).

For DISK devices, estimated capacity is the capacity of all volumes in the storage pool, including any volumes that are varied offline.

For sequential-access storage pools, estimated capacity is the total estimated space of all the sequential-access volumes in the storage pool, regardless of their access mode. At least one volume must be used in a sequential-access storage pool (either a scratch volume or a private volume) to calculate estimated capacity.

For tape and FILE devices, the estimated capacity for the storage pool includes the following factors:

- The capacity of all the scratch volumes that the storage pool already acquired or can acquire. The number of scratch volumes is defined by the MAXSCRATCH parameter on the DEFINE STGPOOL or UPDATE STGPOOL command.
- The total number of available scratch volumes in the tape library.
- Estimated capacity is the smaller number between the MAXSCRATCH value and the total number of available scratch volumes in the tape library.

The calculations for estimated capacity depend on the available space of the storage for the device that is assigned to the storage pool. For FILE storage pools, the capacity for the storage pool is reduced if the available storage is less than the total estimated space of all the FILE volumes in the storage pool. The value that is displayed for capacity is reduced by the size of a FILE volume incrementally as the available space continues to decline.

For Centera, value represents the total capacity of the Centera storage device that is being queried.

Space Trigger Util

Utilization of the storage pool, as calculated by the storage pool space trigger, if any, for this storage pool. You can define space triggers for storage pools that are associated with DISK or FILE device types only.

For sequential access devices, space trigger utilization is expressed as follows as a percentage of the number of used bytes on each sequential access volume relative to the size of the volume and estimated capacity of all existing volumes in the storage pool. It does not include potential scratch volumes. Unlike the calculation for percent utilization, the calculation for space trigger utilization favors creation of new private file volumes by the space trigger over usage of more scratch volumes.

For disk devices, space trigger utilization is expressed as a percentage of the estimated capacity, including cached data. However, it excludes data that is on any volumes that are varied offline. The value for space trigger utilization can be higher than the value for percent migration if you issue QUERY STGPOOL while a file creation is in progress. The value for space trigger utilization is determined by the amount of space that is allocated while the transaction is in progress. The value for percent migration represents only the space that is occupied by committed files. At the end of the transaction, these values are synchronized.

The value for space trigger utilization includes cached data on disk volumes. Therefore, when cache is enabled and migration occurs, the value remains the same because the migrated data remains on the volume as cached data. The value decreases only when the cached data expires or when the space that cached files occupy must be used for noncached files.

Pct Util

An estimate of the utilization of the storage pool, as a percentage.

For sequential access devices, it is a percentage of the number of active bytes on each sequential access volume and the estimated capacity of all volumes in the storage pool. The percentage includes the number of potential scratch volumes that might be allocated.

For disk devices, it is a percentage of the estimated capacity, including cached data and data that is on any volumes that are varied offline. The value for Pct Util can be higher than the value for Pct Migr if you issue this command while a file creation transaction is in progress. The value for Pct Util is determined by the amount of space that is allocated, while the transaction is in progress. The value for Pct Migr represents only the space that is occupied by committed files. At the end of the transaction, these values become synchronized.

The Pct Util value includes cached data on disk volumes. Therefore, when cache is enabled and migration occurs, the Pct Util value remains the same because the migrated data remains on the volume as cached data. The Pct Util value decreases only when the cached data expires or when the space that cached files occupy must be used for noncached files.

For Centera, this represents an estimate of the utilization of the entire Centera storage device, not the storage pool that is being queried.

Pct Migr (primary storage pools only)

An estimate of the percentage of data in the storage pool that can be migrated. The server uses this value and the high and low migration thresholds to determine when to start and stop migration.

For random-access disk devices, this value is specified as a percentage of the value for the estimated capacity, excluding cached data, but including data on any volumes varied offline.

For sequential-access disk devices, this value is specified as a percentage of the value for the estimated capacity. The value includes the capacity of all the scratch volumes that are specified for the pool. For other types of sequential-access

devices, this value is the percentage of the total number of volumes in the pool that contain at least one byte of active data. The total number of volumes includes the maximum number of scratch volumes.

The Pct Util value includes cached data on a volume; the Pct Migr value excludes cached data. Therefore, when cache is enabled and migration occurs, the Pct Migr value decreases but the Pct Util value remains the same because the migrated data remains on the volume as cached data. The Pct Util value decreases only when the cached data expires or when the space that cached files occupy must be used for noncached files.

Pct Logical

The logical occupancy of the storage pool as a percentage of the total occupancy. Logical occupancy is space that is occupied by client files that might or might not be part of an aggregate. A Pct Logical value less than 100% indicates that there is vacant space within aggregates in the storage pool.

High Mig Pct (primary storage pools only)

The high migration threshold, which specifies when the server can begin migration for the storage pool. The server starts migration processes when capacity utilization reaches this threshold.

Low Mig Pct (primary storage pools only)

The low migration threshold, which specifies when the server can stop migration for the storage pool. The server stops migration processes when capacity utilization reaches this threshold.

Migration Delay (primary storage pools only)

The minimum number of days that a file must remain in a storage pool before the server can migrate the file to the next storage pool. For a disk storage pool, the days are counted from the time that the file was stored in the storage pool or last retrieved by a client. For a sequential access storage pool, the days are counted from the time that the file was stored in the storage pool.

Migration Continue (primary storage pools only)

Whether the server continues to migrate files to the next storage pool even if the files have not been in the pool for the number of days that are specified by the migration delay.

Migration Processes

The number of parallel processes that are used for migrating files from a random or sequential access primary storage pool.

Reclamation Processes

The number of parallel processes that are used for reclaiming the volumes in a sequential access primary or copy storage pool.

Next Storage Pool (primary storage pools only)

The storage pool that is the destination for data that is migrated from this storage pool.

Reclaim Storage Pool (primary, sequential access storage pools only)

If specified, the storage pool that is the destination for data that is moved from volumes during reclamation processing. If no pool is specified, by default reclamation processing moves data only among volumes within the same storage pool.

Maximum Size Threshold (primary storage pools only)

The maximum size of files that might be stored in the storage pool.

Access

The access mode for data in the storage pool. The following access modes are possible:

Read/Write

The data can be accessed in read-write mode.

Read only

The data can be accessed in read-only mode.

Converting

The storage pool is being converted to a directory-container storage pool.

Conversion Stopped

The process of converting the storage pool to a directory-container storage pool is stopped.

Conversion Cleanup Needed

To convert the storage pool successfully, you must clean up the storage pool. Conversion could not complete because of damaged data. Issue the QUERY CLEANUP command to identify damaged files.

Converted

The storage pool is converted to a directory-container storage pool.

Description

The description of the storage pool.

Overflow Location (sequential access storage pools only)

The location where volumes in the storage pool are stored when they are ejected from an automated library with the MOVE MEDIA command.

Cache Migrated Files? (random access storage pools only)

Whether caching is enabled for files that are migrated to the next storage pool.

Collocate? (sequential access storage pools only)

Whether collocation is disabled or enabled. If collocation is disabled, the value of this field is No. If collocation is enabled, the possible values are Group, Node, and File space.

Reclamation Threshold (sequential access storage pools only)
The threshold that determines when volumes in a storage pool are reclaimed. The server compares the percentage of reclaimable space on a volume to this value to determine whether reclamation is necessary.

Offsite Reclamation Limit
The number of offsite volumes that have space that is reclaimed during reclamation for this storage pool. This field applies only when POOLTYPE=COPY.

Maximum Scratch Volumes Allowed (sequential access storage pools only)
The maximum number of scratch volumes that the server can request for the storage pool.

Number of Scratch Volumes Used (sequential access storage pools only)
The number of scratch volumes that are used in the storage pool.

Delay Period for Container Reuse (container storage pools only)
The number of days that must elapse after all files are deleted from a container before the server reuses the container.

Migration in Progress? (primary storage pools only)
Whether at least one migration process is active for the storage pool.

Amount Migrated (MB) (primary storage pools only)
The amount of data, in megabytes, that is migrated, if migration is in progress. If migration is not in progress, this value indicates the amount of data that was migrated during the last migration. When multiple, parallel migration processes are used for the storage pool, this value indicates the total amount of data that is migrated by all processes.

Elapsed Migration Time (seconds) (primary storage pools only)
The amount of time that elapsed since migration began, if migration is active. If migration is not active, this value indicates the amount of time that is required to complete the last migration. When multiple, parallel migration processes are used for the storage pool, this value indicates the total time from the beginning of the first process until the completion of the last process.

Reclamation in Progress? (sequential access storage pools only)
Whether a reclamation process is active for the storage pool.

Last Update by (administrator)
The name of the administrator that is defined or most recently updated the storage pool.

Last Update Date/Time
The date and time that an administrator defined or most recently updated the storage pool.

Storage Pool Data Format
The type of data format that is used to write data to this storage pool (for example NATIVE, NETAPPDUMP, CELERRADUMP, or NDMPDUMP).

Copy Storage Pool(s)
The copy storage pools that are listed have data that is simultaneously written to them when data is backed up or archived to the primary storage pool queried by this command.

Active Data Pool(s)
The active-data pools that are listed here have data that is simultaneously written to them when data is backed up to the primary storage pool queried by this command.

Continue Copy on Error?
Whether a server continues to write data to other copy storage pools in the list or ends the entire transaction when a write failure occurs to one of the copy pools in the list. This field applies only to primary random-access and primary sequential-access storage pools.

CRC Data
Whether data is validated by a cyclic redundancy check (CRC) when data is transferred during data storage and retrieval on a device.

Reclamation Type
Whether volumes in this storage pool are reclaimed by threshold or by SnapLock retention date.

Overwrite Data when Deleted
The number of times data will be physically overwritten after it is deleted from the database.

Deduplicate Data?
Whether data in the storage pool is deduplicated.

Processes for Identifying Duplicates
The number of duplicate-identification processes that are specified as the default for the storage pool. The number of duplicate-identification processes that are specified in this field might not equal the number of duplicate-identification processes that are running.

Compressed
Whether the storage pool is compressed.

Additional space for protected data

The amount of space, in MB, that is used to protect data from remote servers. This is the total amount of space used for data received from other servers as a result of running the PROTECT STGPOOL command.

After the PROTECT STGPOOL command is run, the data is not assigned to a node. However, if you run node replication on some or all nodes, then the data is assigned to the nodes and is no longer assigned to the additional space for protected data.

If you do not run node replication, then the data received (after the PROTECT STGPOOL command is run) remains assigned to the additional space for protected data.

Total Unused Pending Space

The amount of space that is scheduled to become available in a directory-container storage pool. The space is occupied by deduplicated data extents that will be removed from the storage pool when the time period specified by the REUSEDDELAY parameter on the DEFINE STGPOOL command expires.

Deduplication Savings

The amount and percentage of data that is saved in the storage pool by using data deduplication.

Compression Savings

The amount of data that is saved in the storage pool by compression.

Total Space Saved

The total amount of data that was saved in the storage pool.

Auto-copy Mode

Indicates whether data is written simultaneously to copy storage pools or active-data pools during client store sessions, server import processes, server data migration processes, or all three operations. The value CLIENT indicates either client store or server import operations. The value ALL indicates that simultaneous-write operations occur whenever this pool is a target for any of the eligible operations.

If the storage pool is a copy storage pool or an active-data pool or if the simultaneous-write function is disabled, this field is blank.

Contains Data Deduplicated by Client?

Indicates whether the storage pool contains data that was deduplicated by clients. Storage pools that contain data that is deduplicated by clients are not accessible for LAN-free data movement by storage agents V6.1 or earlier.

Tip: This field is blank for container storage pools. You cannot use container storage pools for LAN-free data movement.

Maximum Simultaneous Writers

The maximum number of I/O that can run concurrently on the storage pool.

Protect Processes

The set of protect processes.

Protection Storage Pool

The name of the container storage pool where the data is protected to on the target replication server.

Protect Local Storage Pool(s)

Indicates whether local storage pools are protected.

Reclamation Volume Limit

For container-copy storage pools, indicates the maximum number of volumes that the server reclaims during storage pool protection.

Date of Last Protection to Remote Pool

The date that the storage pool was last protected to a storage pool on a remote server.

Date of Last Protection to Local Pool

The date that the storage pool was last protected to a storage pool on the local server.

Deduplicate Requires Backup?

Indicates whether the sequential storage pool must be backed up if the storage pool contains deduplicated data.

Encrypted

For directory-container or cloud-container storage pools, indicates whether client data is encrypted before it is written to the storage pool.

Pct Encrypted

The percentage of deduplicated client data that is encrypted in the directory-container or cloud-container storage pool.

Cloud Space Allocated (MB)

For cloud storage pools, the amount of space that is allocated to cloud storage, in megabytes.

Cloud Space Utilized (MB)

For cloud storage pools, the space that is used by the cloud storage, in megabytes.

Bucket Name

For cloud storage pools that use Simple Storage Service (S3), the name IBM Spectrum Protect™ assigns to the S3 bucket or IBM® Cloud Object Storage vault. This value can also be the name that you assigned to the bucket by using the BUCKETNAME parameter in the DEFINE STGPOOL command or the UPDATE STGPOOL command.

Local Estimated Capacity

For cloud storage pools that use local storage, the estimated capacity of the local storage in megabytes (M) or gigabytes (G).

Local Pct Util

For cloud storage pools that use local storage, an estimate of the utilization of the local storage component of the cloud storage pool, as a percentage.

Local Pct Logical

For cloud storage pools that use local storage, the logical occupancy of the cloud storage pool as a percentage of the total occupancy. Logical occupancy is space that is occupied by client files that might or might not be part of an aggregate. A Local Pct Logical value less than 100% indicates that there is vacant space within aggregates in the cloud storage pool.

Related commands

Table 1. Commands related to QUERY STGPOOL

Command	Description
CONVERT STGPOOL	Convert a storage pool to a directory-container storage pool.
COPY ACTIVE DATA	Copies active backup data.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE STGPOOL	Deletes a storage pool from server storage.
QUERY STGPOOLDIRECTORY	Displays information about storage pool directories.
UPDATE STGPOOL	Changes the attributes of a storage pool.

AIX Linux Windows

QUERY STGPOOLDIRECTORY (Query a storage pool directory)

Use this command to display information about one or more storage pool directories.

Privilege class

Any administrator can issue this command.

Syntax

```
.-*-----.  
>>-Query STGPOOLDIRECTORY--+-+----->  
    '-directory-'  
  
.-ACCESS---Any-----.  
>--+-+----->  
    '-STGpool---pool_name-' '-ACCESS---+READWrite---+'  
                                +-READOnly----+  
                                +-DESTROYED---+  
                                +-Any-----+  
                                '-UNAVAILABLE-'  
  
.-Format---Standard-----.
```

```
>-----<
'-Format--Standard--'
'-Detailed-'
```

Parameters

directory

Specifies the storage pool directory to query. This parameter is optional.

*

Specifies that an asterisk (*) represents a wildcard character. Use wildcard characters such as an asterisk to match any characters. Alternatively, you can use a question mark (?) or a percent sign (%) to match exactly one character. This is the default.

directory

Specifies the storage pool directory. If you do not specify a value for this parameter, all storage pool directories are displayed. The maximum length of the storage pool directory is 1024.

STGpool

Specifies the name of the storage pool to query. If you do not specify a value for this parameter, all storage pool directories are displayed. The maximum length of the storage pool name is 30. This parameter is optional.

ACcESS

Specifies that output is restricted by directory access mode. This parameter is optional. Specify one of the following values:

READWrite

Display all storage pool directories with an access mode of `READWRITE`.

READOnly

Display all storage pool directories with an access mode of `READONLY`.

DESTroyed

Display all storage pool directories with an access mode of `DESTROYED`. The directories are designated as permanently damaged in the storage pool directory.

Any

Display all storage pool directories. This is the default.

UNAVailable

Display directories with an access mode of `UNAVAILABLE`.

Format

Specifies how the information is displayed. This parameter is optional. The default value is `STANDARD`. You can specify one of the following values:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary information for a specific storage pool directory

Display information for the storage pool directory that is named `DPOOL`. See Field descriptions for field descriptions.

```
query stgpoldirectory C:\data
```

Storage Pool Name	Directory	Access
DPOOL	C:\data	Read/Write

Example: Display detailed storage pool directory information

Display details for the storage pool directory named that is named `DPOOL`.

```
query stgpoldirectory stgpool=dpool format=detailed
```

AIX | Linux

```
Storage Pool Name: DPOOL
  Directory: /storage/sampleDir
    Access: Read/Write
  Free Space (MB): 323,170
  Total Space (MB): 476,938
  File System: /storage
  Absolute Path: /storage/data
```

Windows

```
Storage Pool Name: DPOOL
  Directory: /storage2/sampleDir
    Access: Read/Write
  Free Space (MB): 323,170
  Total Space (MB): 476,938
  File System: /storage
  Absolute Path: /storage2/sampleDir
```

Field descriptions

Storage Pool Name

The name of the storage pool.

Directory

The name of the storage pool directory.

Access

The access mode of the data in the storage pool directory.

Free Space (MB)

The amount of space in the storage pool directory, in megabytes, that is not in use.

Total Space (MB)

The total amount of space in the storage pool directory, in megabytes.

File System

The name of the file system where the storage pool directory is located.

Absolute Path

The absolute path name where the storage pool directory is located. The absolute path name contains the name of the root directory and all subdirectories in the path name. All symbolic links are resolved in the absolute path name.

Table 1. Commands related to QUERY STGPOOLDIRECTORY

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
AIX Linux Windows DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
AIX Linux Windows DELETE STGPOOLDIRECTORY	Deletes a storage pool directory from a directory-container or cloud-container storage pool.
AIX Linux Windows UPDATE STGPOOLDIRECTORY	Changes the attributes of a storage pool directory.

QUERY STGRULE (Display storage rule information)

Use this command to display information about storage rules that are defined for storage pools.

Privilege class

Any administrator can issue this command.

Syntax

```

.*----- .-Format----Standard----.
>>-Query STGRULE-----+----->
      '-rule_name-'   '-Format-----Standard--'
                          '-Detailed-'

.-ACTiontype----ANY----- .
>-----+----->
      '-ACTiontype-----ANY-----+'
                          +-AUDit-----+
                          +-GENdedupstats-+
                          +-REClaim-----+
                          '-TIER-----'

.-ACTIVE----ANY----- .
>-----+-----><
      '-ACTIVE-----ANY--+'
                          +-Yes-+
                          '-No--'

```

Parameters

rule_name

Specifies the name of one or more storage rules. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all storage rules are displayed. The maximum length of the name is 30 characters.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. The following values are possible:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

ACTiontype

Specifies the storage action that is completed by the storage rules. The following values are possible:

ANY

All types of storage rules are displayed.

AUDit

Storage rules for audit operations are displayed.

GENdedupstats

Storage rules for data deduplication statistics are displayed.

REClaim

Storage rules for reclaiming cloud-container storage pools are displayed.

TIER

Storage rules for tiering are displayed.

ACTIVE

Specifies whether active or inactive storage rules are displayed. This parameter is optional. The default is ANY. The following values are possible:

ANY

Specifies that all storage rules are displayed.

Yes

Specifies that only active storage rules are displayed.

No

Specifies that only inactive storage rules are displayed.

Example: List all storage rules for all storage pools

Tip: In the output examples, some fields are blank because the item does not apply in the specified environment. Query all storage rules for all storage pools. See QUERY STGRULE (Display storage rule information) for field descriptions.

```
query stgrule
```

Storage Rule Name	Target Storage Pool	Action Type	Active	Source Storage Pools
STGACTION1	CLOUD	Tier	Yes	DIRPOOL1

Example: Display detailed information about a storage rule for tiering

Query detailed information about a storage rule for tiering. See QUERY STGRULE (Display storage rule information) for field descriptions.

```
query stgrule format=detailed
```

```
Storage Rule Name: RULE1
Target Storage Pool: CLOUD1
Action Type: Tier
Active: Yes
Maximum Processes: 8
Start Time: 18:00:00
Delay (in days): 30
Duration:
Description:
Audit Type:
Audit Level:
Node Name:
Filespace names:
Name Type:
Code Type:
Percent Unused:
Last Exe Date/Time:
Source Storage Pools: DIRPOOL1
```

Example: Display detailed information about a storage rule for auditing storage pools

Query detailed information about a storage rule for auditing storage pools. See QUERY STGRULE (Display storage rule information) for field descriptions.

```
query stgrule format=detailed
```

```
Storage Rule Name: AUDIT
Target Storage Pool: CTR
Action Type: Audit
Active: Yes
Maximum Processes: 4
Start Time: 11:42:36
Delay (in days): 7
Duration:
Description:
Audit Type: Extent
Audit Level: 5
Node Name:
Filespace names:
Name Type:
Code Type:
Percent Unused:
Last Exe Date/Time: 01/19/2018 11:43:31
Source Storage Pools:
```

Example: Display detailed information about a storage rule for generating data deduplication statistics

Query detailed information about a storage rule for generating data deduplication statistics. See QUERY STGRULE (Display storage rule information) for field descriptions.

```
query stgrule format=detailed
```

```
Storage Rule Name: GEN1
Target Storage Pool: DIRPOOL
Action Type: GenDedupStats
Active: Yes
Maximum Processes: 8
```

```
      Start Time: 12:06:46
    Delay (in days): 1
      Duration:
    Description:
    Audit Type:
    Audit Level:
      Node Name: *
    Filespace names: *
      Name Type: SERVER
      Code Type: BOTH
    Last Exe Date/Time: 01/18/2018 12:07:10
    Source Storage Pools:
```

Example: Display detailed information about a storage rule for reclaiming space in cloud-container storage pools

Query detailed information about a storage rule for reclaiming space in cloud-container storage pools. See QUERY STGRULE (Display storage rule information) for field descriptions.

```
query stgrule format=detailed
```

```
      Storage Rule Name: RECLAIM
    Target Storage Pool: CLOUD1
      Action Type: Reclaim
        Active: Yes
    Maximum Processes: 8
      Start Time: 9:04:16
    Delay (in days):
      Duration: 120
    Description:
    Audit Type:
    Audit Level:
      Node Name: *
    Filespace names: *
      Name Type:
      Code Type:
    Percent Unused: 50
    Last Exe Date/Time: 01/30/2018 12:07:10
    Source Storage Pools:
```

Field descriptions

Storage Rule Name

The name of the storage rule.

Target Storage Pool

The name of the target storage pool.

Action Type

The type of storage rule.

Active

Indication of whether the storage rule is active or inactive.

Maximum Processes

The number of maximum processes per storage pool.

Tip: For tiering storage rules, this value specifies the maximum number of processes for the source storage pool. For audit storage rules, you cannot set a maximum process value. The server automatically sets and adjusts the number of maximum processes during audit operations.

Start Time

The starting time of the window when the storage rule runs.

Delay (in days)

The number of days to wait before the storage rule operation occurs. For audit storage rules, the number represents the interval, in days, between audit operations. For tiering storage rules, the number represents the minimum number of days that an object must remain in a source storage pool before it is moved to a target storage pool.

Duration

The number of minutes that the storage rule processes the data when all associated processes are completed. No value indicates that processing continues until complete.

Description

A description of the storage rule.

Audit Type

- The type of audit operation.
- Audit Level
 - The level of audit operation.
- Filespace names
 - The names of one or more affected file spaces.
- Name Type
 - Indication of how the server interprets file space names.
- Code Type
 - Indicates the type of file spaces that are included.
- Percent Unused
 - Specifies the percentage of unused space in reclamation storage rules.
- Last Exe Date/Time
 - Specifies the last date and time when the storage rule was run.
- Source Storage Pools
 - The name of the source storage pool or pools.

Related commands

Table 1. Commands related to QUERY STGRULE

Command	Description
DEFINE STGRULE (auditing)	Defines a storage rule for auditing storage pools.
DEFINE STGRULE (data deduplication statistics)	Defines a storage rule for generating data deduplication statistics.
DEFINE STGRULE (reclaiming)	Defines a storage rule for reclaiming cloud-container storage pools.
DEFINE STGRULE (tiering)	Defines a storage rule for tiering.
DELETE STGRULE	Deletes storage rules.
UPDATE STGRULE (auditing)	Updates a storage rule for auditing storage pools.
UPDATE STGRULE (data deduplication statistics)	Updates a storage rule for generating data deduplication statistics.
UPDATE STGRULE (reclaiming)	Updates a storage rule for reclaiming cloud-container storage pools.
UPDATE STGRULE (tiering)	Updates a tiering storage rule.

QUERY SUBSCRIBER (Display subscriber information)

Use this command on a configuration manager to display information about subscribers and their profile subscriptions.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query SUBSCRIBer--+-----+----->
                        .-*-----
                        '-server_name-'

.-PROFile---*-----
>--+-----+----->>
  '-PROFile---profile_name-'

```

Parameters

server_name

Specifies the name of a managed server for which subscription information is displayed. You can use wildcard characters to specify multiple server names. This parameter is optional. The default is all managed servers.

PROFILE

Specifies a profile name for which information is displayed. You can use wildcard characters to specify multiple profile names. This parameter is optional. The default is all profiles.

Example: List a configuration manager's profile subscriptions

Display subscriber information for all profile subscriptions to this configuration manager. See Field descriptions for field descriptions.

```
query subscriber
```

Subscriber	Profile name	Is current?	Last update date/time
SERVER2	DEFAULT_PROFILE	Yes	Thu, May 14, 1998 01:14:42 PM
SERVER2	SETUP	Yes	Thu, May 14, 1998 01:14:42 PM

Field descriptions

Subscriber

The name of the subscriber (managed server).

Profile name

The name of the profile.

Is current?

Whether the subscription has been refreshed with the current information associated with the profile. Possible values are:

Yes

The managed server is current.

No

The managed server is not current. If this field is NO after the profile has been refreshed, check the server messages for error conditions that might cause the refresh to fail.

Unknown

Either the managed server has a more recent version of the profile than the configuration manager, or the profile no longer exists on the configuration manager, but the subscription is still associated with the profile.

Last update date/time

Specifies the date and time that configuration information for the subscription was successfully distributed to the subscriber.

Related commands

Table 1. Commands related to QUERY SUBSCRIBER

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.

QUERY SUBSCRIPTION (Display subscription information)

Use this command on a managed server to display profile subscription information.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query SUBSCRIPTION-----*-----<<
                              +-----+
                              |'-profile_name-'|
```

Parameters

profile_name

Specifies the name of the profile for which subscription information is displayed. You can use wildcard characters to specify multiple names. This parameter is optional. The default is all profiles.

Example: Display description information

Display subscription information for all profiles.

```
query subscription
```

```
Configuration      Profile name      Last update
manager            -----
-----
SERVER1            ADMIN_INFO        Thu, May 14, 1998
                   01:35:13 PM
SERVER1            DEFAULT_PROFILE   Thu, May 14, 1998
                   01:35:13 PM
SERVER1            EMPLOYEE         Thu, May 14, 1998
                   01:35:13 PM
```

Field descriptions

Configuration manager

The name of the configuration manager.

Profile name

The name of the profile.

Last update date/time

When the most recent configuration information was successfully distributed to the subscriber.

Related commands

Table 1. Commands related to QUERY SUBSCRIPTION

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.

QUERY SYSTEM (Query the system configuration and capacity)

Use this command to obtain consolidated information about the server's configuration and capacity.

This command consolidates output from select statements, SHOW commands, and other IBM Spectrum Protect™ commands. Output is generated from several IBM Spectrum Protect commands, for example:

- QUERY ASSOCIATION
- QUERY COPYGROUP

- QUERY DATAMOVER
- QUERY DB
- QUERY DBSPACE
- QUERY DEVCLASS
- QUERY DIRSPACE
- QUERY DOMAIN
- QUERY LIBRARY
- QUERY LOG
- QUERY MGMTCLASS
- QUERY OPTION
- QUERY PROCESS
- QUERY REPLRULE
- QUERY SCHEDULE
- QUERY SERVER
- QUERY SESSION
- QUERY STATUS
- QUERY STGPOOL
- QUERY VOLHISTORY
- QUERY VOLUME

Privilege class

Any administrator can issue this command.

Syntax

```
>>-Query SYStem-----<<
```

Example: View consolidated system information

Issue the QUERY SYSTEM command to obtain consolidated system information. For sample outputs for these query commands, see the individual commands.

```
query system
```

Related commands

Table 1. Commands related to QUERY SYSTEM

Command	Description
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.
QUERY DEVCLASS	Displays information about device classes.
QUERY DOMAIN	Displays information about policy domains.
QUERY LOG	Displays information about the recovery log.
QUERY MGMTCLASS	Displays information about management classes.
QUERY OPTION	Displays information about server options.
QUERY PROCESS	Displays information about background processes.
QUERY SCHEDULE	Displays information about schedules.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
QUERY STGPOOL	Displays information about storage pools.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
QUERY VOLUME	Displays information about storage pool volumes.

QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)

Use this command to display the status of the SET TAPEALERTMSG command. You can enable or disable tape alerts. When enabled, IBM Spectrum Protect™ can retrieve diagnostic information from a tape or library device and display it using ANR messages. When disabled, IBM Spectrum Protect will not query a device for this information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-Query TAPEAlertmsg-----<<
```

Example: Display the status of the QUERY TAPEALERTMSG command

Use the QUERY TAPEALERTMSG command to determine if tape alerts are to be retrieved from devices and displayed in the form of ANR messages.

```
query tapealertmsg
```

```
ANR2017I Administrator SERVER_CONSOLE issued command:
QUERY TAPEALERTMSG
ANR8960I QUERY TAPEALERTMSG: The display of Tape Alerts from SCSI
devices is Enabled.
```

Related commands

Table 1. Commands related to QUERY TAPEALERTMSG

Command	Description
SET TAPEALERTMSG	Specifies whether tape and library devices report diagnostic information to the server.

QUERY TOC (Display table of contents for a backup image)

Use this command to display directory and file information contained in the table of contents (TOC) for a specified backup image. This command does not load table of contents information into the IBM Spectrum Protect™ database. The specified table of contents are read from a storage pool each time the QUERY TOC command is issued.

This command cannot be issued from the server console. If the table of contents is stored on removable media, a mount point is required and output is delayed while the storage pool volume is mounted.

Privilege class

To issue this command you must have either system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax

```
>>-Query TOC--node_name--filesystem_name----->
>--+-----+----->
  '-CREATIONDate----date--CREATIONTime----time-'
  .-Format----Standard----.
>--+-----+----->>
  '-Format----+Standard+-'
      '-Detailed-'
```

Parameters

node_name (Required)

Specifies the name of the NAS node to which the table of contents (TOC) belongs. You cannot use wildcards to specify this name.

filesystem_name (Required)

Specifies the name of the file space to which the table of contents belongs. The file space name you specify cannot contain wildcard characters.

CREATIONDate

Specifies the creation date of the backup image for which the table of contents is to be displayed. This parameter is optional. If you specify CREATIONDATE, you must also specify CREATIONTIME. If you do not specify these parameters, the contents of the latest backup image for the specified node and file space will be displayed, provided that this image has a table of contents. You can only specify the creation date as the following:

Value	Description	Example
MM/DD/YYYY	A specific date	05/15/2002

This specifies that you want to display the contents of the backup image created on this date. You can obtain this date from the output of the QUERY NASBACKUP command.

CREATIONTime

Specifies the creation time of the backup image for which the table of contents is to be displayed. This parameter is optional. If you specify CREATIONTIME, you must also specify CREATIONDATE. If you do not specify these parameters, the contents of the latest backup image for the specified node and file space will be displayed, provided that this image has a table of contents. You can only specify the creation time as the following:

Value	Description	Example
HH:MM:SS	A specific time on the specified creation date.	10:30:08

This specifies that you want to display the contents of the backup image created on this time for the specified date. You can obtain this time from the output of the QUERY NASBACKUP command.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the files.

Detailed

Specifies that complete information is displayed for the files, including the hexadecimal representation of each file or directory name.

Example: Display detailed table of contents information for a specific node

Use the QUERY TOC command to display information in the table of contents belonging to NAS node NETAPP in the file space /vol/vol1 created on 12/06/2002 at 11:22:46. Specify a detailed format.

```
query toc netapp /vol/vol1 creationdate=12/06/2002 creationtime=11:22:46
format=detailed
```

```
Objects in the image backed up on 12/06/2002 11:22:46
for filesystem /vol/vol1 in node NETAPP:
```

```
Object Name: /.etc
```

```

Hexadecimal Object Name: 2f657463
Object Type: Directory
Object Size: 4,096
Last data Modification Date/Time: 07/31/2002 14:21:19

Object Name: /.etc/oldmaps/ndmp
Hexadecimal Object Name: 2f6574632f6f6c646d6170
732f6e646d70
Object Type: Directory
Object Size: 4,096
Last data Modification Date/Time: 07/31/2002 14:21:19

Object Name: /.etc/oldmaps/ndmp/TSM
/vol/vol1/3df0e8fd
Hexadecimal Object Name: 2f6574632f6f6c646d6170
732f6e646d702f54534d2
02f766f6c2f766f6c312f3
364663065386664
Object Type: File
Object Size: 36,864
Last data Modification Date/Time: 12/06/2002 11:14:22

```

Field descriptions

Object Name
The name of the object.

Hexadecimal Object Name
The name of the object in hexadecimal format.

Object Type
The type of the object.

Object Size
The size of the object.

Last data Modification Date/Time
The date and time the object was last modified.

Related commands

Table 1. Commands related to QUERY TOC

Command	Description
BACKUP NODE	Backs up a network-attached storage (NAS) node.
QUERY NASBACKUP	Displays information about NAS backup images.
RESTORE NODE	Restores a network-attached storage (NAS) node.

QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)

Use this command to query a virtual file space mapping definition.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query VIRTUALFSmapping ----->
.
.*-----
>+-----+----->>
|
|'.*-----|
|'-node_name--+-----+'
|'-virtual_filespace_name-'

```

Parameters

node_name

Specifies the client node to which the virtual file space belongs. You can use wildcard characters to specify this name. This parameter is optional. The default is all client node names. You must specify a value for this parameter if you specify a virtual file space name.

virtual_file_space_name

Specifies the name of the virtual file space mappings to be queried. You can use wildcard characters to specify this name. This parameter is optional. If a value is not specified, all virtual file space mappings are queried. Virtual file space mapping names are case sensitive. Use the QUERY VIRTUALFSMAPPING command to determine the correct capitalization for the virtual file space mapping to be queried.

Example: Display virtual file spaces for a specific node

Display the currently defined virtual file spaces for node NAS1. See Field descriptions for field descriptions.

```
query virtualfsmapping nas1
```

Node Name	Virtual Filespace Mapping Name	Filespace Name	Path	Hexadecimal Path?
NAS1	/mikesdir	/vol/vol2	/mikes	No
NAS1	/tmpdir	/vol/vol1	/tmp	No
NAS1	/nonASCIIIDir	/vol/vol3	2f73657276657231	Yes

Field descriptions

Node Name

Specifies the name of the client node.

Virtual Filespace Mapping Name

Specifies the name of the virtual file space mapping.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Path

Specifies the path to the client node.

Hexadecimal Path

Indicates whether the path is hexadecimal.

Related commands

Table 1. Commands related to QUERY VIRTUALFSMAPPING

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

QUERY VOLHISTORY (Display sequential volume history information)

Use this command to display sequential volume history information. To save sequential volume history information to one or more files, use the BACKUP VOLHISTORY command.

Use the VOLUMEHISTORY server option to specify one or more volume history files. After the server is restarted, IBM Spectrum Protect™ updates volume information in both the database and the files.

Use the QUERY BACKUPSET command to query specified backup set information.

Privilege class

Any administrator can issue this command.

Syntax

```

>>-Query VOLHistory--+-BEGINDate---earliest_date-.----->
                        '-BEGINDate---date-----'

      .-ENDDate---current_date-.  .-BEGINTime---00:00:00-.
>--+-----+-----+-----+-----+-----+----->
      '-ENDDate---date-----'  '-BEGINTime---time-----'

      .-ENDTime---current_time-.  .-Type---All------.
>--+-----+-----+-----+-----+-----+-----><
      '-ENDTime---time-----'  '-Type---+All-----+'
                                +-BACKUPSET---+
                                +-DBBackup---+
                                +-DBRpf-----+
                                +-DBSnapshot--+
                                +-EXPort-----+
                                |         (1)  |
                                +-REmote-----+
                                +-RPFile-----+
                                +-RPFSnapshot--+
                                +-STGDelete---+
                                +-STGNew-----+
                                '-STGReuse----'

```

Notes:

1. This parameter is only available on AIX, HP-UX, Linux, Solaris and Windows operating systems.

Parameters

BEGINDate

Specifies that you want to display information beginning with records created on the specified date. This parameter is optional. The default is the earliest date for which history information exists.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified	TODAY-7 or -7. To display information beginning with records created a week ago, specify BEGINDATE=TODAY-7 or BEGINDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.

Value	Description	Example
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies that you want to display information ending with records created on the specified date. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY-1 or -1. To display records created up to yesterday, specify ENDDATE=TODAY-1 or ENDDATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies that you want to display information beginning with records created at the specified time. This parameter is optional. The default is midnight (00:00:00).

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	12:33:28
NOW	The current time on the specified begin date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 with BEGINTIME=NOW+03:00 or BEGINTIME=+03:00, IBM Spectrum Protect displays records with a time of 12:00 or later on the begin date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue this command at 9:00 with BEGINTIME=NOW-03:30 or BEGINTIME=-03:30, IBM Spectrum Protect displays records with a time of 5:30 or later on the begin date.

ENDTime

Specifies that you want to display information ending with records created at the specified time on the end date. This parameter is optional. The default is the current time.

You can specify the time using one of the values below:

Value	Description	Example
-------	-------------	---------

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+03:00 or ENDTIME=+03:00, IBM Spectrum Protect displays records with a time of 12:00 or later on the end date.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30 If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME=-3:30, IBM Spectrum Protect displays records with a time of 5:30 or earlier on the end date.

Type

Specifies the type of records to display from the volume history file. This parameter is optional. The default is ALL. Possible values are:

All

Specifies all records.

BACKUPSET

Specifies to display only information about backup set volumes.

DBBackup

Specifies to display only records that contain information about full and incremental database backup volumes, that is with the volume types of BACKUPFULL and BACKUPINCR.

DBRpf

Specifies to display only records that contain information about full and incremental database backup volumes and recovery plan file object volumes (volume types of BACKUPFULL, BACKUPINCR, and RPFIL).

DBSnapshot

Specifies to display only records that contain information about volumes used for database snapshot backups.

EXPort

Specifies only records that contain information about export volumes.

REMOte

Specifies to display only records that contain information about volumes used by library clients.

RPFil

Specifies to display only records that contain information about file objects of a recovery plan that are saved on a target server and that were created assuming database full and incremental backups. The parameter displays only records about recovery plan files that are saved on another IBM Spectrum Protect server by using the server-to-server virtual volume function for IBM Spectrum Protect.

RPFSnapshot

Specifies to display only records that contain information about file objects of a recovery plan that are saved on a target server and that were created assuming database snapshot backups. RPFSnapshot only displays records about recovery plan files that are saved on another IBM Spectrum Protect server by using the server-to-server virtual volume function for IBM Spectrum Protect.

STGDelete

Specifies only records that contain information about deleted sequential storage pool volumes.

STGNew

Specifies only records that contain information about new sequential access storage volumes.

STGReuse

Specifies only records that contain information about reused sequential storage pool volumes.

Example: Display volume history information for a storage pool volume

Display volume history information for a storage pool volume stored in the database. See Field descriptions for field descriptions. Issue the command:

```
query volhistory type=stgnew
```

```

Date/Time: 02/25/2011 18:28:06
Volume Type: STGNEW
Backup Series:
Backup Operation:
Volume Seq:
Device Class: FILE
Volume Name: /adsmfct/server/prv011
Volume Location:
Command:
Database Backup ID High:
Database Backup ID LOW:
Database Backup Home Position:
Database Backup HLA:
Database Backup LLA:
Database Backup Total Data Bytes (MB):
Database Backup total Log Bytes (MB):
Database Backup Block Num High:
Database Backup Block Num Low:
Database Backup Stream Id:
Database Backup Volume Sequence for Stream:

```

Note: The volume history file will contain additional fields that do not appear in the query output. These fields are specific to database backup and restore support. They are not intended for use or modification by IBM Spectrum Protect administrators. The fields will be bracketed with a message indicating these are for IBM Spectrum Protect internal use only and not meant to be modified.

Example: Display volume history information for a database backup volume

Display volume history information for a database backup volume stored in the database. See Field descriptions for field descriptions. Issue the command:

```
query volhistory type=dbb
```

```

Date/Time: 02/25/2011 18:28:06
Volume Type: BACKUPFULL
Backup Series: 176
Backup Operation: 0
Volume Seq: 0
Device Class: FILE
Volume Name: /adsmfct/server/prv011
Volume Location:
Command:
Database Backup ID High: 0
Database Backup ID LOW: 0
Database Backup Home Position: 0
Database Backup HLA:
Database Backup LLA:
Database Backup Total Data Bytes (MB): 0
Database Backup total Log Bytes (MB): 0
Database Backup Block Num High: 0
Database Backup Block Num Low: 0
Database Backup Stream Id: 1
Database Backup Volume Sequence for Stream: 10,001

```

Note: The volume history file will contain additional fields that do not appear in the query output. These fields are specific to database backup and restore support. They are not intended for use or modification by IBM Spectrum Protect administrators. The fields will be bracketed with a message indicating these are for IBM Spectrum Protect internal use only and not meant to be modified.

Field descriptions

Date/Time

The date and time that the volume was created.

Volume Type

The type of volume:

BACKUPFULL

Full database backup volume.

BACKUPINCR	Incremental database backup volume.
BACKUPSET	Client backup set volume.
DBSNAPSHOT	Snapshot database backup volume.
EXPORT	Export volume.
REMOTE	A volume used on the library client, which is the IBM Spectrum Protect server named in the Volume Location field. See the volume history on the server that is the library client to get details about how the volume is used.
RPFIL	Recovery plan file object volume created assuming full and incremental database backups.
RPFSnapshot	Recovery plan file object volume created assuming snapshot database backups.
STGDELETE	Deleted sequential access storage pool volume.
STGNEW	Added sequential access storage pool volume.
STGREUSE	Reused sequential access storage pool volume.

Backup Series

The value of this field depends on the volume type:

- For BACKUPFULL or BACKUPINCR volume types: the backup series identifier.
- For the DBSNAPSHOT volume type: the identifier of the backup series that is associated with the DBSNAPSHOT entry.
- For the RPFIL volume type: the identifier of the backup series that is associated with the RPFIL entry.
- For the RPFSnapshot volume type: the identifier of the backup series that is associated with the RPFSnapshot entry.
- For BACKUPSET volume types: this field is blank.
- For all other volume types: always 0.

A backup series is a full backup and all incremental backups that apply to that full backup. Another series begins with the next full backup of the database.

Backup Operation

For BACKUPFULL or BACKUPINCR volume types: the operation number of this backup volume within the backup series. The full backup within a backup series is operation 0. The first incremental backup for that full backup is operation 1, the second incremental backup is operation 2, and so on.

For DBSNAPSHOT volume types: the operation number of this DBSNAPSHOT volume within the DBSNAPSHOT series.

For all other volume types: always 0.

This field is blank when the volume type is BACKUPSET.

Volume Seq

The sequence or position of the volume within the backup series.

- For BACKUPFULL or BACKUPINCR volume types: the sequence, or position, of the volume within the backup series. Volume sequence 1 identifies the first volume used for the first operation (a full backup), and so on. For example, if the full backup occupies three volumes, these volumes are identified as volume sequence 1, 2, and 3, respectively. The first volume of the next operation (the first incremental backup) is then volume sequence 4.
- For BACKUPSET volume types: the sequence, or position, of the volume within the BACKUPSET series.
- For DBSNAPSHOT volume types: the sequence, or position, of the volume within the DBSNAPSHOT series. Volume sequence 1 identifies the first volume used for the first DBSNAPSHOT operation, and so on.
- For EXPORT volume types: the sequence number of the volume when it was used for exporting data.
- For RPFIL volume types: the value of this field is always one (1).
- For all other volume types: always 0.

Device Class

The name of the device class associated with this volume.

Volume Name

The name of the volume.

Volume Location

The location of the volume. This information is available only for the following volume types:

- BACKUPFULL
- BACKUPINCR
- EXPORT
- REMOTE
- RPPFILE

For the volume type of REMOTE, this location field is the server name of the library client that owns this volume.

For the volume type of RPPFILE, this location field is the server name defined in the device class definition used by the PREPARE command when the DEVCLASS parameter is specified.

Command

When the volume type is EXPORT or BACKUPSET and the volume sequence is 1 (for example, the first volume), this field shows the command that was used to generate the volume. If the EXPORT or BACKUPSET is on more than one volume, the command is displayed with the first volume but not with any of the other volumes.

For any volume type other than EXPORT or BACKUPSET, this field is blank.

Tip: The following fields are not used by IBM Spectrum Protect servers that are V6.3 or later. However, the fields are displayed for compatibility with earlier releases.

- Database Backup ID High
- Database Backup ID Low
- Database Backup Home Position
- Database Backup HLA
- Database Backup LLA
- Database Backup Total Data Bytes (MB)
- Database Backup Total Log Bytes (MB)
- Database Backup Block Num High
- Database Backup Block Num Low

Related commands

Table 1. Commands related to QUERY VOLHISTORY

Command	Description
BACKUP VOLHISTORY	Records volume history information in external files.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
PREPARE	Creates a recovery plan file.
QUERY RPPFILE	Displays information about recovery plan files.
QUERY BACKUPSET	Displays backup sets.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.

QUERY VOLUME (Query storage pool volumes)

Use this command to display information about one or more storage pool volumes.

Privilege class

Any administrator can issue this command.

Syntax

```

.*-----
>>-Query Volume-----+----->
      '-volume_name-'

>+-----+----->
|          .-,----- . |
|          V              | |
| '-ACCess-----+READWrite-----+-' |
|               +READOnly-----+ |
|               +UNAVailable--+ |
|               +OFFsite-----+ |
|               '-DESTROYed---' |

.*-----
>+-----+-----+-----+----->
|          .-,----- . | '-STGpool-----pool_name-'
|          V              | |
| '-STatus-----+ONline--+-' |
|               +OFFline-+ |
|               +EMPTy---+ |
|               +PENding-+ |
|               +FILLing-+ |
|               '-FULL----' |

.-DEVclass-----*-----
>+-----+----->
| '-DEVclass-----device_class_name-'

.-Format-----Standard-----
>+-----+-----><
| '-Format-----+Standard-+-'
|               '-Detailed-'

```

Parameters

volume_name

Specifies the volume to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a name, all storage pool volumes are included in the query.

ACCess

Specifies that output is restricted by volume access mode. This parameter is optional. You can specify multiple access modes by separating the modes with commas and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by access mode. Possible values are:

READWrite

Display volumes with an access mode of READWRITE. Client nodes and server processes can read from and write to files stored on the volumes.

READOnly

Display volumes with an access mode of READONLY. Client nodes and server processes can read only files that are stored on the volumes.

UNAVailable

Display volumes with an access mode of UNAVAILABLE. Client nodes and server processes cannot access files that are stored on the volumes.

OFFsite

Display copy storage pool volumes with an access mode of OFFSITE. The volumes are at offsite locations from which they cannot be mounted.

DESTROYed

Display primary storage pool volumes with an access mode of DESTROYED. The volumes are designated as permanently damaged.

Status

Specifies that output is restricted by volume status. This parameter is optional. You can specify multiple status values by separating values with commas and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by volume status. Possible values are:

ONline

Display random access volumes that are available to the server.

Offline

Display random access volumes that are not available to the server.

EMPTy

Display sequential access volumes that have no data.

PENding

Display volumes with a status of PENDING. These volumes might be sequential-access volumes from which all files were deleted, but for which the time specified by the REUSEDELAY parameter on the DEFINE STGPOOL command has not elapsed. These volumes might also be random-access disk volumes that were deleted, but that still contain discarded data that is waiting to be shredded. After the data is shredded, the volume will be physically deleted.

FILLing

Display sequential access volumes that the server has written to but has not yet filled to capacity.

FULL

Display sequential access volumes that the server filled.

STGPool

Specifies the storage pool to include in the query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a storage pool name, all storage pools are included in the query.

DEVclass

Specifies the device class to include in the query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a device class name, all devices are included in the query.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

AIX

Linux

Example: List all file storage pool volumes

Display information on all storage pool volumes with the device class name of FILE. See Field descriptions for field descriptions.

```
query volume devclass=file
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Volume Status
/FCT/SERVER/COV011	COPYSTG	FILE	0.0 M	0.0	Pending
/FCT/SERVER/COV012	COPYSTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/COV013	COPYSTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/PRV011	PRIMESTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/PRV012	PRIMESTG	FILE	0.0 M	0.0	Empty

Windows

Example: List all storage pool volumes with the same prefix

Display information on all storage pool volumes that are prefixed with the name ATF. See Field descriptions for field descriptions.

```
query volume atf*
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Volume Status
ATF001	8MMPPOOL	8MMTAPE	4.8 G	18.2	Filling
ATF002	8MMPPOOL	8MMTAPE	4.8 G	18.2	Filling

AIX

Linux

Example: Display detailed information about a specific storage pool volume

Display details about the storage pool volume named /fct/server/cov011. See Field descriptions for field descriptions.

```
query volume cov011 format=detailed
```

```

        Volume Name: /FCT/SERVER/COV011
        Storage Pool Name: COPYSTG
        Device Class Name: DISK
        Estimated Capacity: 10.0 M
        Scaled Capacity Applied:
            Pct Util: 6.7
            Volume Status: On-line
            Access: Read/Write
        Pct. Reclaimable Space: 3.2
        Scratch Volume?: Yes
        In Error State?: No
        Number of Writable Sides: 1
        Number of Times Mounted: 11
        Write Pass Number: 1
        Approx. Date Last Written: 04/14/1998 16:17:26
        Approx. Date Last Read: 04/01/1998 13:26:18
        Date Became Pending:
        Number of Write Errors: 0
        Number of Read Errors: 0
        Volume Location:
        Volume is MVS Lanfree Capable: No
        Last Update by (administrator): COLLIN
        Last Update Date/Time: 05/01/1998 14:07:27
        Begin Reclaim Period:
        End Reclaim Period:
        Logical Block Protected:
        Drive Encryption Key Manager:

```

Windows

Example: Display detailed information about a specific storage pool volume

Display details about the storage pool volume WPDV00. See Field descriptions for field descriptions.

```

query volume wpdv00 format=detailed

        Volume Name: WPDV00
        Storage Pool Name: TAPEPOOL
        Device Class Name: TAPE
        Estimated Capacity: 5.8 M
        Scaled Capacity Applied:
            Pct Util: 0.1
            Volume Status: On-line
            Access: Read/Write
        Pct. Reclaimable Space: 3.2
        Scratch Volume?: Yes
        In Error State?: No
        Number of Writable Sides: 1
        Number of Times Mounted: 11
        Write Pass Number: 1
        Approx. Date Last Written: 04/14/1998 16:17:26
        Approx. Date Last Read: 04/01/1998 13:26:18
        Date Became Pending:
        Number of Write Errors: 0
        Number of Read Errors: 0
        Volume Location:
        Volume is MVS Lanfree Capable: No
        Last Update by (administrator): COLLIN
        Last Update Date/Time: 05/01/1998 14:07:27
        Begin Reclaim Period:
        End Reclaim Period:
        Logical Block Protected:
        Drive Encryption Key Manager:

```

Example: Display detailed information about a storage pool volume with a specific device class

Display details about a volume in a storage pool with a device class name of FILECLASS. See Field descriptions for field descriptions.

```

query volume devclass=fileclass format=detailed

```

Windows	Volume Name: Z:\WORM_CFS\0000000E.BFS	
AIX	Linux	Volume Name: /WORM_FILESYS/0000000E.BFS

```

Storage Pool Name: FILEPOOL
Device Class Name: FILECLASS
Estimated Capacity: 2.0 G
Scaled Capacity Applied:
  Pct Util: 0.0
  Volume Status: Filling
  Access: Read/Write
Pct. Reclaimable Space: 0.0
  Scratch Volume?: Yes
  In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 1
  Write Pass Number: 1
Approx. Date Last Written: 03/22/2004 15:23:46
Approx. Date Last Read: 03/22/2004 15:23:46
  Date Became Pending:
Number of Write Errors: 0
Number of Read Errors: 0
  Volume Location:
Volume is MVS Lanfree Capable: No
Last Update by (administrator):
  Last Update Date/Time: 03/22/2004 15:23:46
  Begin Reclaim Period: 03/22/2005
  End Reclaim Period: 04/22/2005
Logical Block Protected:
Drive Encryption Key Manager:

```

Example: Display detailed information about a specific storage pool volume

Display details about a storage pool volume that is named 000642. The volume is in a storage pool that is associated with a 3592 device class. See Field descriptions for field descriptions.

```

query volume 000642 format=detailed

      Volume Name: 000642
      Storage Pool Name: 3592POOL
      Device Class Name: 3592CLASS
      Estimated Capacity: 2.0 G
Scaled Capacity Applied:
  Pct Util: 0.0
  Volume Status: Filling
  Access: Read/Write
Pct. Reclaimable Space: 0.0
  Scratch Volume?: Yes
  In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 1
  Write Pass Number: 1
Approx. Date Last Written: 03/22/2004 15:23:46
Approx. Date Last Read: 03/22/2004 15:23:46
  Date Became Pending:
Number of Write Errors: 0
Number of Read Errors: 0
  Volume Location:
Volume is MVS Lanfree Capable: No
Last Update by (administrator):
  Last Update Date/Time: 03/22/2004 15:23:46
  Begin Reclaim Period: 03/22/2005
  End Reclaim Period: 04/22/2005
Logical Block Protected: Yes
Drive Encryption Key Manager: IBM Spectrum Protect

```

Field descriptions

Volume Name

The name of the storage pool volume.

Storage Pool Name

The storage pool to which the volume is defined.

Device Class Name

The device class that is assigned to the storage pool.

Estimated Capacity

The estimated capacity of the volume, in megabytes (M), gigabytes (G), or terabytes (T).

For DISK devices, this value is the capacity of the volume.

For sequential access devices, this value is an estimate of the total space available on the volume, which is based on the device class.

Scaled Capacity Applied

The percentage of capacity to which a volume is scaled. For example, a value of 20 for a volume whose maximum capacity is 300 GB indicates that the volume can store only 20 percent of 300 GB, or 60 GB. This attribute applies only to IBM® 3592 devices.

Pct Util

An estimate of the utilization of the volume. The utilization includes all space that is occupied by both files and aggregates, including empty space within aggregates.

For DISK volumes, the utilization also includes space that is occupied by cached data.

Volume Status

The status of the volume.

Access

Whether the volume is available to the server.

Pct. Reclaimable Space (sequential access volumes only)

The amount of space on this volume that can be reclaimed because data has expired or been deleted. This value is compared to the reclamation threshold for the storage pool to determine whether reclamation is necessary. Reclaimable space includes empty space within aggregates.

When determining which volumes in a storage pool to reclaim, the server first determines the reclamation threshold. The reclamation threshold is indicated by the value of the THRESHOLD parameter on the RECLAIM STGPOOL command or, if that value was not specified, the value of the RECLAIM parameter in a storage pool definition. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose that storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70.

For volumes that belong to a SnapLock storage pool, the value is displayed but is not used.

Scratch Volume? (sequential access volumes only)

Whether this volume is returned to scratch when the volume becomes empty.

In Error State?

Whether the volume is in an error state. The server cannot write to volumes in an error state.

Number of Writable Sides

This information is reserved for IBM Spectrum Protect™.

Number of Times Mounted

The number of times that the server opened the volume for use. The number of times that the server opened the volume is not always the same as the number of times that the volume was physically mounted in a drive. After a volume is physically mounted, the server can open the same volume multiple times for different operations, for example for different client backup sessions.

Write Pass Number (sequential access volumes only)

The number of times the volume was written to from the beginning to the end.

Approx. Date Last Written

The approximate date on which the volume was last written.

Approx. Date Last Read

The approximate date on which the volume was last read.

Date Became Pending

The date that the status of the volume was changed to pending.

Number of Write Errors

The number of writing errors that occurred on the volume.

Number of Read Errors

The number of reading errors that occurred on the volume.

Volume Location

The location of the volume.

Volume is MVS Lanfree Capable

Whether the volume is LAN-free capable. A LAN-free capable volume is one that was defined and used (at least once) by the IBM Spectrum Protect z/OS® data manager server.

Last Update by (administrator)

The administrator that defined or most recently updated the volume.

Last Update Date/Time

When the volume was defined or most recently updated.

Begin Reclaim Period

Represents the date after which the server begins reclaiming this volume, but not later than the date represented by the end reclaim period. If, when the reclaim period begins, there are files on the volume that have not expired, they are moved to a new WORM volume during reclamation processing. This field displays a date only if this volume is in a storage pool for which the value of the RECLAMATIONTYPE parameter is SNAPLOCK.

If more than one archive is stored on the same volume, the start of the volume's reclamation period is based on the date of the most recent archive. For SnapLock volumes, the RETVer parameter of the DEFINE COPYGROUP command determines how long an archive is stored. If RETVer is set to 100 days, the volume's reclamation period will start 100 days after the first archive is stored on it. If a second archive is stored on the same volume, the reclamation start date will be adjusted to 100 days after the new archive is stored. If the RETVer value is changed after the first archive is stored, the latest reclamation date will apply for all of the archives on the volume. For example, assume RETVer is set to 100 for an initial archive, but is then changed to 50. If a second archive is stored on the volume three days after the first, the reclamation period will not start until 100 days after the first archive was stored.

End Reclaim Period

Represents the date by which the IBM Spectrum Protect must complete reclamation processing on this volume to ensure continued protection of the data. It also represents the Last Access Date physical file attribute in the NetApp Filer, which prevents the file from being deleted until after that date. This field displays a date only if this volume is in a storage pool for which the value of the RECLAMATIONTYPE parameter is SNAPLOCK.

Drive Encryption Key Manager

The drive encryption key manager. This field applies only to volumes in a storage pool that is associated with a device type of 3592, LTO, or ECARTRIDGE.

Logical Block Protected

Specifies whether logical block protection is enabled for the volume. You can use logical block protection only with the following types of drives and media:

- IBM LTO5 and later
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later
- Oracle StorageTek T10000C and T10000D drives

Related commands

Table 1. Commands related to QUERY VOLUME

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE VOLUME	Updates the attributes of storage pool volumes.
VARY	Specifies whether a disk volume is available to the server for use.

QUIT (End the interactive mode of the administrative client)

Use this command to end an administrative client session in interactive mode.

You cannot use the QUIT command from the SERVER_CONSOLE administrative ID, or the console, batch, or mount modes of the administrative client.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-QUIT-----<<
```

Parameters

None.

Example: End an interactive administrative client session

End an administrative client session in the interactive mode.

```
quit
```

Related commands

None.

RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)

Use this command to reclaim volumes in a sequential-access storage pool. Reclamation does not move inactive versions of backup data from volumes in active-data pools.

This command cannot be used for the following types of storage pools:

- Container-copy storage pools. Space in these storage pools is reclaimed as part of the processing that is done by PROTECT STGPOOL commands.
- Storage pools with one of the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
- Storage pools that use a CENTERA device class.
- Storage pools that use a Write Once Read Many (WORM) device class. Reclamation is not necessary because WORM volumes are not reusable, but you can run reclamation to consolidate data onto fewer volumes.

Use this command only if you are not going to use automatic reclamation for the storage pool. This command accepts the values of the RECLAIMPROCESS and RECLAIMSTGPOOL attributes of the storage pool definition. This command also accepts the values of the OFFSITERECLAIMLIMIT and RECLAIM parameters of the storage pool definition, if not overridden by the OFFSITERECLAIMLIMIT and THRESHOLD command parameters.

Tips:

- When you issue this command, duplicate data in a primary storage pool, copy storage pool, or active-data pool that is set up for data deduplication is removed.
- When you use this command to restore deduplicated objects to the same storage pool, any duplicate data blocks are replaced with references to deduplicated extents.

For storage pools defined with RECLAMATIONTYPE=SNAPLOCK, this command also deletes empty WORM FILE volumes that exceeded their reclaim period.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool that is being reclaimed and the reclaim storage pool, if applicable.

Syntax

```
>>-RECLaim STGpool--pool_name--+-+-----+----->
                                     '-Threshold---number-'
                                     .-Wait---No-----
>--+-----+----->
   '-Duration---minutes-' '-Wait---+No--+'
                                     '-Yes-'

>--+-----+-----><
   '-OFFSITERECLAIMLimit---number_of_volumes-'
```

Parameters

pool_name (Required)

Specifies the storage pool in which volumes are to be reclaimed.

DURation

Specifies the maximum number of minutes that the reclamation runs before it is automatically canceled. You can specify a number 1 - 9999. This parameter is optional.

After the specified number of minutes elapses, the next time the server checks the reclamation process the server stops the reclamation process. The server checks the reclamation process when the server mounts another eligible volume from the storage pool that is being reclaimed. The server also checks the reclamation process when the server begins to reclaim a new batch of files from the currently mounted volume. As a result, the reclamation can run longer than the value you specified for this parameter.

Until the server checks the reclamation process, there is no indication the duration period expired. When the server stops the reclamation process, the server issues message ANR4927W: Reclamation terminated for volume xxx - duration exceeded.

If you do not specify this parameter, the process stops only when no more volumes meet the threshold.

If you specify a duration value for reclamation of a copy storage pool with offsite volumes, you might cause the reclamation to end before any volumes are reclaimed. In most situations when you initiate reclamation for a copy storage pool with offsite volumes, consider limiting the number of offsite volumes to be reclaimed rather than limiting the duration. For details, see the OFFSITERECLAIMLIMIT parameter.

THreshold

Specifies the percentage of reclaimable space on a volume that makes it eligible for reclamation. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the server database. Reclaimable space also includes unused space.

You can specify a number 1 - 99. This parameter is optional. If not specified, the RECLAIM attribute of the storage pool definition is used.

To determine the percentage of reclaimable space for a volume, issue the QUERY VOLUME command and specify FORMAT=DETAILED. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined into a single target volume.

OFFSITERECLAIMLimit

Specifies the maximum number of offsite storage pool volumes that the server tries to reclaim. This parameter is valid only for copy storage pools. You can specify a number 0 - 99999. This parameter is optional. If not specified, the OFFSITERECLAIMLIMIT attribute of the storage pool definition is used.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. You can specify one of the following values:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is processed. Messages that are created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

If you cancel this process, some files might already be moved to new volumes before the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. Output messages are displayed to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Reclaim volumes in a sequential-access storage pool

Reclaim volumes in the storage pool named TAPEPOOL. Specify that reclamation ends as soon as possible after 60 minutes.

```
reclaim stgpool tapepool duration=60
```

Related commands

Table 1. Commands related to RECLAIM STGPOOL

Command	Description
CANCEL PROCESS	Cancels a background server process.
MIGRATE STGPOOL	Migrates files from a primary storage pool to the next storage pool in the hierarchy.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.

RECONCILE VOLUMES (Reconcile differences in the virtual volume definitions)

Issue this command from the source server to reconcile differences between virtual volume definitions on the source server and archive files on the target server. IBM Spectrum Protect™ finds all volumes of the specified device class on the source server and all corresponding archive files on the target server. The target server inventory is also compared to the local definition for virtual volumes to see if inconsistencies exist.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>-REConcile Volumes-----+----->
                             .-*-----+
                             +-----+
                             '-device_class_name-'

.-Fix---No-----
>--+-----+-----<
  '-Fix-----No---'
```

'-Yes-'

Parameters

device_class_name

Specifies the device class name of the virtual volumes. If you do not specify a name, IBM Spectrum Protect reconciles all virtual volumes. This parameter is optional.

FIX

Specifies whether or not IBM Spectrum Protect attempts to correct any identified inconsistencies. This parameter is optional. The default is NO. Possible values are:

No

Specifies that IBM Spectrum Protect does not fix any inconsistencies.

Yes

Specifies that IBM Spectrum Protect makes the following corrections:

- IBM Spectrum Protect marks as unavailable storage pool volumes on the source server that cannot be located on the target server. Volumes that are only found in the volume history, such as database backups and import and export volumes, are reported as being inconsistent.
- Archive files on the target server that do not correspond to any virtual volumes on the source server are marked for deletion from the target server.

The following table shows the details of the actions taken:

FIX=	At the Source Server	At the Target Server	Action
NO	Volumes exist	No files exist	Report error
		Files exist but are marked for deletion	
		Active files exist but attributes do not match	
	Volumes do not exist	Active files exist	Report error
	Files exist but are marked for deletion	None	
YES	Volumes exist	No files exist	Report error Storage pool volumes: Marked as unavailable
		Files exist but marked for deletion	Report error Storage pool volumes: If attributes match, mark files on the target server as active again, mark volumes on the source server as unavailable, and recommend that an AUDIT VOLUME be done to verify the data. If attributes do not match, mark volumes as unavailable.
		Active files exist but attributes do not match	Report error Storage pool volumes: Mark as unavailable and recommend that an AUDIT VOLUME be done to verify the data.
	Volumes do not exist	Active files exist	Mark files for deletion on the target server.
		Files exist but marked for deletion	None

Example: Reconcile differences in the virtual volume definitions

Reconcile the differences between all virtual volumes definitions on the source server and archive files on the target server to correct any inconsistencies.

```
reconcile volumes remotel fix=yes
```

Related commands

Table 1. Commands related to RECONCILE VOLUMES

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE SERVER	Deletes the definition of a server.
QUERY SERVER	Displays information about servers.
UPDATE SERVER	Updates information about a server.

REGISTER commands

Use the REGISTER commands to define or add objects to IBM Spectrum Protect™.

- REGISTER ADMIN (Register an administrator ID)
- REGISTER LICENSE (Register a new license)
- REGISTER NODE (Register a node)

REGISTER ADMIN (Register an administrator ID)

Use this command to add an administrator to the server. After registration, the administrator can issue a limited set of commands, including all query commands. To provide additional privileges, use the GRANT AUTHORITY command.

Privilege class

To issue this command, you must have system privilege.

When you register an administrator with the same name as an existing node, be aware of the administrator authentication method and the SSLREQUIRED setting. Any node that has the same name as the administrator that is being registered inherits those settings.

For users of Lightweight Directory Access Protocol (LDAP) servers:

- The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.
- Do not specify an administrative user ID that matches a node name. If the administrative user ID matches the node name, you might see unexpected behavior because of automatic password changes that update the same password twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update operation might fail.

Syntax

```
>>-REGister Admin--admin_name--+-----+----->
                                     '-password-'
>--+-----+-----+-----+----->
| (1) | | '-CONTACT----text-'
|-----PASSExp----days-'
.-FORCEPwreset----No-----
>--+-----+-----+-----+----->
|'-FORCEPwreset----+No--+-'
|                                     '-Yes-'
>--+-----+-----+-----+----->
|'-EMAILAddress----userID@node-'
(2)
```


EMAILAddress

Specifies the email address for this administrator.

AUTHentication

This parameter specifies the authentication method for the administrator user ID. Specify one of the following values: LDAP or LOCAL. The parameter is optional and defaults to LOCAL. The default can change to LDAP if you use the SET DEFAULTAUTHENTICATION command and specify LDAP.

Local

Specifies that the local IBM Spectrum Protect server database is used.

LDap

Specifies that the administrator user ID authenticates passwords with an LDAP directory server. Passwords that authenticate with an LDAP directory server are case-sensitive.

Tip: A password is not required if you register an administrator and select `AUTHENTICATION=LDAP`. At logon, you are prompted for a password.

SSLrequired (deprecated)

Specifies whether the administrator user ID must use the Secure Sockets Layer (SSL) protocol to communicate between the IBM Spectrum Protect server and the backup-archive client. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Important: Beginning with IBM Spectrum Protect Version 8.1.2 software and Tivoli® Storage Manager Version 7.1.8 software, this parameter is deprecated. Validation that was enabled by this parameter is replaced by the `SESSIONSECURITY` parameter, which is enforced by the `SESSIONSECURITY` parameter. The `SSLREQUIRED` parameter is ignored. Update your configuration to use the `SESSIONSECURITY` parameter.

SESSIONSECurity

Specifies whether the administrator must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRict

Specifies that the strictest security settings are enforced for the administrator. The `STRICT` value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the server and the administrator. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option.

To use the `STRICT` value, the following requirements must be met to ensure that the administrator can authenticate with the server:

- Both the administrator and server must be using IBM Spectrum Protect software that supports the `SESSIONSECURITY` parameter.
- The administrator must be configured to use the TLS 1.2 protocol for SSL sessions between the server and the administrator.

Administrators set to `STRICT` that do not meet these requirements are unable to authenticate with the server.

TRANSitional

Specifies that the existing security settings are enforced for the administrator. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the `STRICT` value.

If `SESSIONSECURITY=TRANSITIONAL` and the administrator has never met the requirements for the `STRICT` value, the administrator will continue to authenticate by using the `TRANSITIONAL` value. However, after an administrator meets the requirements for the `STRICT` value, the `SESSIONSECURITY` parameter value automatically updates from `TRANSITIONAL` to `STRICT`. Then, the administrator can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for `STRICT`. In addition, after an administrator successfully authenticates by using a more secure communication protocol, the administrator can no longer authenticate by using a less secure protocol. For example, if an administrator that is not using SSL is updated and successfully authenticates by using TLS 1.2, the administrator can no longer authenticate by using no SSL protocol or TLS 1.1. This restriction also applies when you use functions such as command routing or server-to-server export, when the administrator authenticates to the IBM Spectrum Protect server as an administrator from another server.

ALert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No

Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the QUERY MONITORSETTINGS command.

Example: Register an administrator

Define an administrator, LARRY, with the password PASSWORDONE. You can identify LARRY as second-shift personnel by specifying this information with the CONTACT parameter. Issue the command:

```
register admin larry passwordone contact='second shift'
```

Example: Register an administrator ID and set the authentication method

Define an administrator ID for Harry so that Harry can authenticate to an LDAP server. Issue the command:

```
register admin harry authentication=ldap
```

Example: Register an administrator and enforce strict session security

Register an administrator named Harry, and require Harry to use the strictest security settings to authenticate with the server. Issue the command:

```
register admin harry sessionsecurity=strict
```

Related commands

Table 1. Commands related to REGISTER ADMIN

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
LOCK ADMIN	Prevents an administrator from accessing IBM Spectrum Protect.
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE ADMIN	Removes an administrator from the list of registered administrators.
RENAME ADMIN	Changes an IBM Spectrum Protect administrator's name.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UNLOCK ADMIN	Enables a locked administrator to access IBM Spectrum Protect.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes that are associated with a client node.

Related tasks:

Naming Tivoli Storage Manager objects

Related reference:

[Ssl client option](#)

REGISTER LICENSE (Register a new license)

Use this command to register new licenses for server components, including IBM Spectrum Protect™ (base), IBM Spectrum Protect Extended Edition, and IBM Spectrum Protect for Data Retention.

Licenses are stored in enrollment certificate files. The enrollment certificate files contain licensing information for the server product. The NODELOCK file preserves the licensing information for your installation. Your license agreement determines what you are licensed to use, even if you cannot use the REGISTER LICENSE command to register all components. You are expected to comply with the license agreement and use only what you have purchased. Use of the REGISTER LICENSE command implies that you agree to and accept the license terms specified in your license agreement.

Important:

- Before upgrading from a previous version of IBM Spectrum Protect, you must delete or rename the NODELOCK file.
- To unregister licenses, you must erase the NODELOCK file in the server instance directory of your installation, and reregister any previously registered licenses.
- You cannot register licenses for IBM Spectrum Protect for Mail, IBM Spectrum Protect for Databases, IBM Spectrum Protect for ERP, and IBM Spectrum Protect for Space Management.

To generate a report that can help you understand the license requirements for your system, run the QUERY PVUESTIMATE command. The report contains estimates of the number of client devices and PVU totals for server devices. The estimates are not legally binding.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REGister LICense--FILE--==--+tsmbasic.lic+-----><
                                +-tsmee.lic----+
                                +-dataret.lic--+
                                '/*.lic-----'
```

Parameters

FILE

Specifies the name of the enrollment certificate file containing the license to be registered. The specification can contain a wildcard (*). Enter the complete file name or a wildcard in place of the file name. The file names are case-sensitive. The following values can be used:

tsmbasic.lic

To license base IBM Spectrum Protect.

tsmee.lic

To license IBM Spectrum Protect Extended Edition. This includes the disaster recovery manager, large libraries, and NDMP.

dataret.lic

To license IBM Spectrum Protect for Data Retention. This is required to enable Data Retention Protection as well as Expiration and Deletion Suspension (Deletion Hold).

*.lic

To license all IBM Spectrum Protect licenses for server components.

Example: Register a license

Register the base IBM Spectrum Protect license.

```
register license file=tsmbasic.lic
```

Related commands

Table 1. Commands related to REGISTER LICENSE

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY LICENSE	Displays information about licenses and audits.
QUERY PVUESTIMATE	Displays processor value unit estimates.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.

REGISTER NODE (Register a node)

Use this command to register a node to the server.

This command can create an administrative user ID with client owner authority over the node. You can use this administrative user ID to access the web backup-archive client from remote locations through a web browser.

Tip:

- In earlier product releases, the REGISTER NODE command automatically created an administrative user ID whose name matched the node name. Beginning with IBM Spectrum Protect™ V8.1, the REGISTER NODE command does not automatically create an administrative user ID that matches the node name.
- If you plan to use the LAN-free option with this node, you must register an administrative ID that matches the node name. To register the administrative ID, use the USERID parameter or manually register the administrator and grant owner authority to the node.

If a client requires a different policy domain than STANDARD, you must register the client node with this command or update the registered node.

Requirement: When you set `sslrequired=serveronly` in a REGISTER NODE command, the admin SSLREQUIRED setting reverts to YES. To use a non-SSL session with a storage agent, rename the admin with the identical name by issuing the RENAME ADMIN command.

For users of Lightweight Directory Access Protocol (LDAP) servers: The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.

When you register or update a node, you can specify whether damaged files on the node can be recovered from a replication server. Files can be recovered only if all the following conditions are met:

- Version 7.1.1 or later, is installed on the source and target replication servers.
- The REPLRECOVERDAMAGED system parameter is set to ON. The system parameter can be set by using the SET REPLRECOVERDAMAGED command.
- The source server includes at least one file that is marked as damaged in the node that is being replicated.
- The node data was replicated before the damage occurred.

The following table describes how parameter settings affect the recovery of damaged, replicated files.

Table 1. Settings that affect the recovery of damaged files

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result
OFF	YES, NO, or not specified	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
OFF	ONLY	YES or NO	An error message is displayed because files cannot be recovered when the REPLRECOVERDAMAGED system parameter is set to OFF.

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result
ON	YES	YES or NO	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	NO	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
ON	ONLY	YES or NO	Damaged files are recovered from the target replication server, but standard node replication does not occur.
ON	Not specified	YES	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	Not specified	NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-REGister Node--node_name--+-----+----->
                                     '-password-'
                                     .-Userid---NONE-----
>+-----+-----+-----+----->
| (1) | | '-Userid---NONE---+' |
|-----PASSExp---days- | | '-user_id-' |
                                     .-Domain---STANDARD-----
>+-----+-----+-----+----->
'-CONtact---text-' '-Domain-----domain_name---'
                                     .-COMPression---Client----- .-ARCHDElete---Yes-----
>+-----+-----+-----+----->
'-COMPression---+Client+-' '-ARCHDElete---+Yes+-'
                                     +-Yes-----+-No--
                                     '-No-----'
                                     .-BACKDElete---No-----
>+-----+-----+-----+----->
'-BACKDElete---+No---+'
                                     '-Yes-'
>+-----+-----+-----+----->
'-CLOptset---option_set_name-'
                                     .-FORCEPwreset---No----- .-Type---Client-----
>+-----+-----+-----+----->
'-FORCEPwreset---+No---+' '-Type---+Client---+'
                                     '-Yes-' | (2) |
                                               +-NAS-----+
```

```

                                '-Server--'
>--+-----+-----+-----+-----+-----+----->
  '-URL-----url-'  '-UTILITYurl-----utility_url-'
. -MAXNUMMP-----1----- . -AUTOFSRename-----No----- .
>--+-----+-----+-----+-----+-----+----->
  '-MAXNUMMP-----number-'  '-AUTOFSRename-----+Yes-----+-'
                                   +-No-----+
                                   '-Client-'

. -KEEPMP-----No----- . (3)
>--+-----+-----+-----+-----+-----+----->
  '-KEEPMP-----+No--+-'
                                   '-Yes-'

. -VALIDateprotocol-----No----- .
>--+-----+-----+-----+-----+-----+----->
  '-VALIDateprotocol-----+No-----+-'
                                   +-Dataonly+
                                   '-All-----'

. -TXNGroupmax-----0----- .
>--+-----+-----+-----+-----+-----+----->
  '-TXNGroupmax-----+0-----+-'
                                   '-number-'

. -DATAWritepath-----ANY----- .
>--+-----+-----+-----+-----+-----+----->
  '-DATAWritepath-----+ANY-----+-'
                                   +-LAN-----+
                                   '-LANFree-'

. -DATAReadpath-----ANY----- .
>--+-----+-----+-----+-----+-----+----->
  '-DATAReadpath-----+ANY-----+-'
                                   +-LAN-----+
                                   '-LANFree-'

>--+-----+-----+-----+-----+-----+----->
  '-TARGETLevel-----V.R.M.F-'

. -SESSIONINITiation-----Clientorserver----- .
>--+-----+-----+-----+-----+-----+-----+----->
  '-SESSIONINITiation-----+Clientorserver-----+-'
                                   '-SERVEROnly--HLAddress-----ip_address--LLAddress-----tcp_port-'

>--+-----+-----+-----+-----+-----+----->
  '-HLAddress-----ip_address--LLAddress-----tcp_port-'

>--+-----+-----+-----+-----+-----+----->
  '-EMAILAddress-----userID@node-'

. -DEDUPlication-----Clientorserver----- .
>--+-----+-----+-----+-----+-----+----->
  '-DEDUPlication-----+Clientorserver+-'
                                   '-SERVEROnly-----'

. -BACKUPINITiation-----All----- .
>--+-----+-----+-----+-----+-----+----->
  |                                     (4) |
  '-BACKUPINITiation-----+All--+-----'
                                   '-ROOT-'

>--+-----+-----+-----+-----+-----+----->
  '-REPLState-----+Enabled--+-'
                                   '-DISabled-'

. -BKREPLRuledefault-----DEFAULT----- .
>--+-----+-----+-----+-----+-----+-----+----->
  | (5) |
  '------BKREPLRuledefault-----+ALL_DATA-----+-'
                                   +-ACTIVE_DATA-----+
                                   +-ALL_DATA_HIGH_PRIORITY-----+

```

```

++ACTIVE_DATA_HIGH_PRIORITY++
+-DEFAULT-----+
'-NONE-----'

.-ARREPLRuledefault----DEFAULT-----
>-----+-----+----->
| (5) |
'|-----ARREPLRuledefault----+--ALL_DATA-----+|
|                                     +-ALL_DATA_HIGH_PRIORITY-+|
|                                     +-DEFAULT-----+|
|                                     '-NONE-----'|

.-SPREPLRuledefault----DEFAULT-----
>-----+-----+----->
| (5) |
'|-----SPREPLRuledefault----+--ALL_DATA-----+|
|                                     +-ALL_DATA_HIGH_PRIORITY-+|
|                                     +-DEFAULT-----+|
|                                     '-NONE-----'|

.-RECOVERDamaged----Yes-----
>-----+-----+----->
'|RECOVERDamaged----+-Yes-+|
|                               '-No--'|

.-ROLEOVERRIDE----Userreported-----
>-----+-----+----->
'|ROLEOVERRIDE----+Client-----+|
|                               +-Server-----+|
|                               +-Other-----+|
|                               '-Userreported-|

(6)
.------AUTHentication----Local-.
>-----+-----+----->
'|AUTHentication----+-Local-+---|
|                               '-LDap--'|

(7)
.-SSLrequired----Default-----
>-----+-----+----->
'|SSLrequired----+-Yes-----+|
|                               +-No-----+|
|                               +-Default----+|
|                               '-SERVERonly-|

.-SESSIONSECurity----TRANSitional-----
>-----+-----+----->
'|SESSIONSECurity----+-STRict-----+|
|                               '-TRANSitional-|

.-SPLITLARGEObjects----Yes-----
>-----+-----+----->>
'|SPLITLARGEObjects----+-Yes-+|
|                               '-No--'|

```

Notes:

1. The PASSEXPCOMMAND does not apply to administrators who authenticate with a Lightweight Directory Access Protocol (LDAP) directory server.
2. This parameter is only available for AIX®, Linux, and Windows operating systems.
3. The VALIDATEPROTOCOL parameter is deprecated.
4. The BACKUPINITIATION parameter is ignored if the client node operating system is not supported.
5. You can specify the BKREPLRULEDEFAULT, ARREPLRULEDEFAULT, or SPREPLRULEDEFAULT parameter only if you specify the REPLSTATE parameter.
6. The default value can change if you issued the SET DEFAULTAUTHENTICATION command and specified LDAP.
7. The SSLREQUIRED parameter is deprecated.

Parameters

node_name (Required)

Specifies the name of the client node to be registered. The maximum length of the name is 64 characters.

You cannot specify a node name of NONE.

password

Specifies the client node password. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

If you authenticate passwords locally with the IBM Spectrum Protect server, you must specify a password. The password is not case-sensitive.

If you authenticate passwords with an LDAP server, do not specify a password on the REGISTER NODE command.

PASSExp

Specifies the number of days the password remains valid. You can set the password expiration period 0 - 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the server common-password expiration period is used. The common password expiration period is 90 days unless changed by issuing the SET PASSEXP command.

You can change the password expiration period by using the UPDATE NODE or SET PASSEXP commands. You can issue the SET PASSEXP command to set a common expiration period for all administrators and client nodes. You can also use the command to selectively set password expiration periods. If you selectively set a password expiration period by using the REGISTER NODE command, the UPDATE NODE command, or the SET PASSEXP command, the expiration period is excluded from common password expiration periods that were created by using the SET PASSEXP command.

You can use the RESET PASSEXP command to reset the password expiration period to the common expiration period. The PASSEXP command does not apply to nodes that authenticate with an LDAP server.

USerid

Specifies the administrative user ID with client owner authority. This parameter is optional. You can specify one of the following values:

NONE

Specifies that no administrative user ID is created. This is the default value.

user_id

Specifies that an administrative user ID is created with the specified name. You can use this parameter to grant client owner authority to an existing administrative user ID.

If you register a node that has the same name as an administrator, the administrator authentication method and SSLREQUIRED setting change to match the authentication method of the node. Passwords that are shared between same-named nodes and administrators are kept synchronized during an authentication change.

If you plan to use the LAN-free option with this node, use the USERID parameter to register an administrative ID that matches the node name.

For users of LDAP servers: If you plan to authenticate the node with an LDAP server, keep the default setting (USERID=NONE) or specify an administrative user ID that differs from the node name. If the administrative user ID matches the node name, you might see unexpected behavior because of automatic password changes that update the same password twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update operation might fail.

CONtact

Specifies a text string of information that identifies the node. The parameter is optional. The maximum length of the text string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks.

DOmain

Specifies the name of the policy domain to which the node is assigned. The parameter is optional. If you do not specify a policy domain name, the node is assigned to the default policy domain (STANDARD).

When a source server is registered as a node, it is assigned to a policy domain. Data from the source server is stored in the storage pool that is specified in the archive copy group of the default management class of that domain.

COMPression

Specifies whether the client node compresses its files before it sends these files to the server for backup and archive. The parameter is optional. The default value is CLIENT.

Restriction: This parameter does not apply to nodes with a type of NAS or SERVER.

You can specify one of the following values:

Client

Specifies that the client determines whether to compress files.

Yes

Specifies that the client node compresses its files before it sends these files to the server for backup and archive.

No

Specifies that the client node does not compress its files before it sends these files to the server for backup and archive.

ARCHDElete

Specifies whether the client node can delete its own archive files from the server. The parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that the client node can delete its own archive files from the server.

No

Specifies that the client node cannot delete its own archive files from the server.

BACKDElete

Specifies whether the client node can delete its own backup files from the server. The parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the client node cannot delete its own backup files from the server.

Yes

Specifies that the client node can delete its own backup files from the server.

CLOptset

Specifies the name of the option set to be used by the client. The parameter is optional.

FORCEPwreset

Specifies whether to force a client to change or reset the password. The parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the password expiration period is set by the SET PASSEXP command. The client does not need to change or reset the password while the client is logging on to the server.

Yes

Specifies that the client node password expires at the next logon. The client must change or reset the password then. If a password is not specified, you receive an error message.

Restriction: For nodes that authenticate with an LDAP server, password expiration is set by using LDAP server utilities. For this reason, do not specify FORCEPWRESET=YES if you specify AUTHENTICATION=LDAP.

Type

Specifies the type of node that is being registered. The parameter is optional. The default value is CLIENT. You can specify one of the following values:

Client

Specifies that the client node is a Backup-Archive Client, IBM Spectrum Protect for Space Management client, or application client.

NAS

Specifies that the node is a network-attached storage (NAS) file server whose data is protected by using NDMP operations. The node name cannot be SERVER.

Note: The name of the NAS node must be the same as the data mover. Therefore, the name cannot be changed after a corresponding data mover is defined.

Server

Specifies that the client node is a source server that is being registered on the target server.

URL

Specifies the URL of the IBM Spectrum Protect web client that is configured on the client system. You can use the URL in a web browser and in the Operations Center to remotely manage the client node.

This parameter is optional. The URL must include the DNS name or IP address of the client system, and the port number that is defined on the client system for the IBM Spectrum Protect web client. For example,
`http://client.mycorp.com:1581`

UTILITYUrl

Specifies the address of the IBM Spectrum Protect client management services that are configured on the client system. This URL is used by the Operations Center to access client log files so that you can remotely diagnose client issues from the Operations Center.

This parameter is optional. You can specify a URL of up to 200 characters in length. The URL must start with `https`. It includes the DNS name or IP address of the client system, and the port number that is defined on the client system for the IBM Spectrum Protect client management services. For example, `https://client.mycorp.com:9028`

If you omit the port number, the Operations Center uses the port number 9028, which is the default port number when you install the client management services on the client system.

MAXNUMMP

Specifies the maximum number of mount points a node is allowed to use on the server or storage agent only for operations such as backup, archive, and IBM Spectrum Protect for Space Management migration. The parameter is optional and does not apply to nodes with a type of NAS or SERVER. The default value is 1. You can specify an integer in the range 0 - 999. A value of 0 specifies that a node cannot acquire any mount point for a client data store operation. The MAXNUMMP value is not evaluated or enforced during client data read operations such as restore, retrieve, and IBM Spectrum Protect for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node and might prevent the data store operations from being able to acquire mount points.

For volumes in a storage pool that is associated with the FILE or CENTERA device type, the server can have multiple sessions to read and one process to write to the same volume concurrently. To increase concurrency and provide efficient access for nodes with data in FILE or CENTERA storage pools, increase the value of the MAXNUMMP parameter.

For nodes that store data into primary storage pools with the simultaneous-write function that is enabled, you must adjust the value of the MAXNUMMP parameter to specify the correct number of mount points for each client session. A client session requires one mount point for the primary storage pool and one mount point for each copy storage pool and each active-data pool.

For server-to-server backup, if one server is at a different version than the other server, set the number of mount points on the target server to a value higher than one. Otherwise, you receive an error.

A storage agent independently tracks the number of points that are used during a client session. If a node has a storage agent that is installed, it might exceed the MAXNUMMP value. The MAXNUMMP value might also be exceeded under conditions where the node does not have to wait for a mount point.

Note: The server might preempt a client operation for a higher priority operation and the client might lose a mount point if no other mount points are available.

KEEPMP

Specifies whether the client node keeps the mount point for the entire session. The parameter is optional. The default value is NO. You can specify one of the following values:

Yes

Specifies that the client node must retain the mount point during the entire session. If policy definitions cause data to be stored to a disk storage pool after the data is stored to a sequential access storage pool, any mount points that are held by the session will not be released.

No

Specifies that the client node releases the mount point during the session. If policy definitions cause data to be stored to a disk storage pool after the data is stored to a sequential access storage pool, any mount points that are held by the session will be released.

AUTOFSRename

Specify whether file spaces are automatically renamed when you upgrade the client system to support Unicode or specify whether file spaces are renamed by the client, if needed. The parameter is optional. The default is NO. Setting the parameter to YES enables automatic renaming, which occurs when the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The automatic renaming changes the names of existing backed-up file spaces that are not in Unicode in server storage. Then, the file spaces are backed up in Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect clients by using Windows, Macintosh OS X, and NetWare operating systems.

After the client with support for Unicode is installed, any new file spaces that the client backs up are stored in server storage by using the UTF-8 code page. UTF-8 is a byte-oriented encoding form that is specified by the Unicode Standard.

You can specify one of the following values:

Yes

Existing file spaces are automatically renamed when you upgrade to a client that supports Unicode and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming occurs whether the client uses the graphical user interface, the command line, or the client scheduler.

For example, the server renames a drive as follows:

```
Original name: D_DRIVE  
New name: D_DRIVE_OLD
```

The new name indicates that the file space is stored on the server in a format that is not Unicode.

No

Existing file spaces are not automatically renamed when the client system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup.

Client

The option AUTOFSRENAME in the client's option file determines whether file spaces are renamed.

By default, the client option is set to PROMPT. When the client system upgrades to a client that supports Unicode and the client runs an IBM Spectrum Protect operation with the graphical user interface or the command line, the program displays a one-time prompt to the user about whether to rename file spaces.

When the client scheduler runs an operation, the program does not prompt for a choice about renaming, and does not rename file spaces. Backups of existing file spaces are sent as before (not in Unicode).

VALIDATEprotocol (deprecated)

Specifies whether IBM Spectrum Protect completes a cyclic redundancy check (CRC) to validate the data that is sent between the client and server. The parameter is optional. The default is NO.

Important: Beginning with IBM Spectrum Protect V8.1.2 and Tivoli® Storage Manager Version 7.1.8, this parameter is deprecated. The validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The VALIDATEPROTOCOL parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

TXNGroupmax

Specifies the number of files per transaction commit that are transferred between a client and a server. The parameter is optional. Client performance might be improved by using a larger value for this option.

The default value is 0. Specifying 0 indicates that the node uses the server global value that is set in the server options file. To use a value other than the server global value, specify a value of 4 through 65,000 for this parameter. The node value takes precedence over the server value.

Attention: Increasing the TXNGROUPMAX value increases the recovery log usage. Higher recovery log usage might increase the risk of running out of log space. Evaluate the performance of each node before you change the parameter.

DATAWritepath

Specifies the transfer path that is used when the client sends data to the server, storage agent, or both, during storage operations such as backup or archive. The parameter is optional. The default is ANY.

Note: If a path is unavailable, the node cannot send any data. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation fails.

You can specify one of the following values:

ANY

Specifies that data is sent to the server, storage agent, or both, by any available path. A LAN-free path is used if one is available. If a LAN-free path is unavailable, the data is moved by using the LAN.

LAN

Specifies that data is sent by using the LAN.

LANFree

Specifies that data is sent by using a LAN-free path.

DATAReadpath

Specifies the transfer path that is used when the server, storage agent, or both read data for a client, during operations such as restore or retrieve. The parameter is optional. The default is ANY.

Note: If a path is unavailable, data cannot be read. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation fails. The value for the transfer path also applies to failover connections. If the value is set to LANFree, failover cannot occur for the node on the secondary server.

You can specify one of the following values:

ANY

Specifies that the server, storage agent, or both use any available path to read data. A LAN-free path is used if one is available. If a LAN-free path is unavailable, the data is read by using the LAN.

LAN

Specifies that data is read by using the LAN.

LANFree

Specifies that data is read by using a LAN-free path.

TARGETLevel

Specifies the client deployment package that is targeted for this node. You can substitute an applicable release package for Version.Release.Modification.Fix (V.R.M.F) Level. For example: `TARGETLevel=6.2.0.0`.

You must specify each segment with a number that is applicable to a deployment package. You cannot use an asterisk in any field as a substitution for a valid number. The parameter is optional.

Restriction: The TARGETLEVEL parameter does not apply to nodes with a type of NAS or SERVER.

SESSIONInitiation

Controls whether the server or the client initiates sessions. The default is that the client initiates sessions. The parameter is optional.

Clientorserver

Specifies that the client might initiate sessions with the server by communicating on the TCP/IP port that is defined with the server option TCPPOINT. Server-prompted scheduling might also be used to prompt the client to connect to the server.

SERVEROnly

Specifies that the server does not accept client requests for sessions. All sessions must be initiated by server-prompted scheduling on the port that is defined for the client with the REGISTER or UPDATE NODE commands. You cannot use the client acceptor, dsmscd, to start the scheduler when SESSIONINITIATION is set to SERVERONLY.

HLAddress

Specifies the client IP address that the server contacts to initiate scheduled events. This parameter must be used when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that are previously used by the client to contact the server.

The address can be specified either in numeric or host name format. If a numeric address is used, it is saved without verification by a domain name server. If the address is not correct, it can cause failures when the server attempts to contact the client. Host name format addresses are verified with a domain name server. Verified names are saved and resolved with Domain Name Services when the server contacts the client.

LLAddress

Specifies the client port number on which the client listens for sessions from the server. This parameter must be used when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that are previously used by the client to contact the server.

The value for this parameter must match the value of client option TCPCLIENTPORT. The default value is 1501.

EMAILAddress

This parameter is used for more contact information. The parameter is optional. The information that is specified by this parameter is not acted upon by IBM Spectrum Protect.

DEDUPLICATION

Specifies where data deduplication can occur for this node. The parameter is optional. You can specify one of the following values:

Clientorserver

Specifies that data that is stored by this node can be deduplicated on either the client or the server. This value is the default. For data deduplication to take place on the client, you must also specify a value of YES for the

DEDUPLICATION client option. You can specify this option in the client option file or in the client option set on the IBM Spectrum Protect server.

SERVEROnly

Specifies that data that is stored by this node can be deduplicated on the server only.

BACKUPINITiation

Specifies whether the non-root user ID on the client node can back up files to the server. The parameter is optional. The default value is ALL, indicating that non-root user IDs can back up data to the server. You can select one of the following values:

All

Specifies that non-root user IDs can back up files to the server. ALL is the default if BACKUPINITIATION is not specified.

ROOT

Specifies that the root user ID can back up files to the server. If you are using the V6.4 or later backup-archive client, authorized users have the same privileges as the root user ID.

Restriction: The attribute is ignored by the server if the backup-archive client connects from an operating system other than AIX, Linux, or Mac OS.

Remember: The application programming interface (API) is affected by the BACKUPINITIATION parameter on the server. By default, all API users are allowed to back up data. Setting the parameter to ROOT on an API node is not recommended.

REPLState

Specifies whether data that belongs to the client node is ready to be replicated. This parameter is optional. Specify this parameter only if you are issuing the REGISTER NODE command on a server that is configured to replicate data to a target replication server. If you register a client node on a source replication server and set up replication for the node, do not register the node on the target replication server. The client node is created automatically on the target server the first time that replication occurs.

You can select one of the following values:

Enabled

Specifies that the client node is configured for replication and is ready to replicate. When you specify this parameter, the replication mode in the client node definition on the source replication server is automatically set to SEND. This setting indicates that data that belongs to the client node is sent to a target server during replication.

When replication first occurs for the client node, the replication state of the node on the target replication server is automatically set to ENABLED. The replication mode on the target replication server is set to RECEIVE. This setting indicates that data that belongs to the client node is received from a source replication server. To determine the replication state and mode, issue the QUERY NODE command on a source or a target replication server.

DISabled

Specifies that the node is configured for replication but that replication does not occur until you enable it.

BKREPLRuledefault, ARREPLRuledefault, and SPREPLRuledefault

Specifies the replication rule that applies to a data type if the file space rules for the data type are set to DEFAULT.

Restriction: You can specify the BKREPLRULEDEFAULT, ARREPLRULEDEFAULT, or SPREPLRULEDEFAULT parameter only if you specify the REPLSTATE parameter.

BKREPLRuledefault

Specifies the replication rule for backup data.

ARREPLRuledefault

Specifies the replication rule for archive data.

SPREPLRuledefault

Specifies the replication rule for space-managed data.

If the file space rules for the data type are set to DEFAULT and you do not specify a rule for the BKREPLRULEDEFAULT, ARREPLRULEDEFAULT, or SPREPLRULEDEFAULT parameter, data is replicated according to the server rule for the data type.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

You can specify the following rules:

ALL_DATA

Replicates active and inactive backup data, archive data, or space-managed data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority. This rule is valid only for BKREPLRULEDEFAULT.

Attention:

If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a release version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the `FORCERECONCILE=YES` parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a release version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data, archive data, or space-managed data. Data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority. This rule is valid only for BKREPLRULEDEFAULT.

DEFAULT

Replicates data according to the server replication rule for backup data.

For example, suppose that you want to replicate the archive data in all the file spaces that belongs to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify `ARREPLRULEDEFAULT=DEFAULT`. Ensure that the file space rules for archive data are also set to DEFAULT and that the server rule for archive data is set to ALL_DATA_HIGH_PRIORITY.

Restriction: If a node is configured for replication, the file space rules are set to DEFAULT after the node stores data on the source replication server.

NONE

Data of the specified type is not replicated.

For example, if you do not want to replicate space-managed data that belongs to a client node, specify `SREPLRULEDEFAULT=NONE`

RECOVERDamaged

Specifies whether damaged files can be recovered for this node from a target replication server. The parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that recovery of damaged files from a target replication server is enabled for this node.

No

Specifies that recovery of damaged files from a target replication server is not enabled for this node.

Tip: The value of the RECOVERDAMAGED parameter is only one of several settings that determine whether damaged files are recovered. For information about how to specify the settings, see Settings that affect the recovery of damaged files.

ROLEOVERRIDE

Specifies whether to override the reported role of the client for processor value unit (PVU) estimation reporting. The default is USERREPORTED. The parameter is optional.

The role reported by the client is either client-device (for example, a workstation) or server-device (for example, file/print server, application server, database). By default, the client reports its role that is based on the client type and the operating system. All clients initially report their role as server-device, except for Backup-Archive Clients running Microsoft Windows workstation distributions (Windows Vista) and Macintosh OS X.

Specify one of the following values:

Client

Specifies a client-device.

Server

Specifies a server-device.

Other

Specifies that this node is not to be used for PVU estimation reporting. This value can be useful when multiple nodes are deployed for a physical system (for example, virtual environments, test nodes, retired nodes, and nodes not in production or clustering).

Usereported

Use the reported role that is provided by the client.

AUTHentication

This parameter specifies the password authentication method for the node. Specify one of the following values: LDAP or LOCAL. The parameter is optional and defaults to LOCAL. The default can change to LDAP if you use the SET DEFAULTAUTHENTICATION command and specify LDAP.

Local

Specifies that the local IBM Spectrum Protect server database is used.

LDap

Specifies that the node uses an LDAP server for password authentication.

SSLrequired (deprecated)

Specifies whether the node must use the Secure Sockets Layer (SSL) protocol to communicate with the IBM Spectrum Protect server. The parameter is optional. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Important: Beginning with IBM Spectrum Protect V8.1.2 software and Tivoli Storage Manager V7.1.8 software, this parameter is deprecated. The validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The SSLREQUIRED parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

SESSIONSECurity

Specifies whether the node must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRict

Specifies that the strictest security settings are enforced for the node. The STRICT value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the server and the node. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option.

To use the STRICT value, the following requirements must be met to ensure that the node can authenticate with the server:

- Both the node and server must be using IBM Spectrum Protect software that supports the SESSIONSECURITY parameter.
- The node must be configured to use the TLS 1.2 protocol for SSL sessions between the server and the node.

Nodes set to STRICT that do not meet these requirements are unable to authenticate with the server.

TRANSitional

Specifies that the existing security settings are enforced for the node. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the STRICT value.

If SESSIONSECURITY=TRANSITIONAL and the node has never met the requirements for the STRICT value, the node will continue to authenticate by using the TRANSITIONAL value. However, after a node meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. Then, the node can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT. In addition, after a node successfully authenticates by using a more secure communication protocol, the node can no longer authenticate by using a less secure protocol. For example, if a node that is not using SSL is updated and successfully authenticates by using TLS 1.2, the node can no longer authenticate by using no SSL protocol or by using TLS 1.1. This restriction also applies when you use functions such as virtual volumes, when the node authenticates to the IBM Spectrum Protect server as a node from another server.

SPLITLARGEObjects

Specifies whether large objects that are stored by this node are automatically split into smaller pieces, by the server, to optimize server processing. The parameter is optional. Specifying Yes causes the server to split large objects (over 10 GB)

into smaller pieces when stored by a client node. Specifying No bypasses this process. Specify No only if your primary concern is maximizing throughput of backups directly to tape. The default value is Yes.

Example: Register a client node that only the root user can back up

Register the client node `mete0rite` with password `KingK0ng` to back up files from only the root user to the server.

```
register node mete0rite KingK0ng
backupinit=root
```

Example: Register a client node and password and set compression on

Register the client node `JOEOS2` with the password `SECRETCODE` and assign this node to the `DOM1` policy domain. This node can delete its own backup and archive files from the server. All files are compressed by the client node before they are sent to the server. This command automatically creates a `JOEOS2` administrative user ID with password `SECRETCODE`. In addition, the administrator now has client owner authority to the `JOEOS2` node.

```
register node joeos2 secretcode domain=dom1
archdelete=yes backdelete=yes
compression=yes
```

Example: Grant client owner authority for an existing administrative user

Grant client owner authority to an existing administrative user ID, `HELPAADMIN`, when you register the client node `JAN`. This step would not automatically create an administrator ID named `JAN`, but would grant client owner authority for this node to the `HELPAADMIN` administrator.

```
register node jan pwd1safe userid=helpadmin
```

Example: Register a NAS file server node that uses NDMP operations

Register a node name of `NAS1` for a NAS file server that is using NDMP operations. Assign this node to a special NAS domain.

```
register node nas1 pwd4nas1 domain=nasdom type=nas
```

Example: Register a node and specify the maximum number of files per transaction commit

Register a node name of `ED` and set the `TXNGROUPMAX` to 1000.

```
register node ed pw459twx txngroupmax=1000
```

Example: Register a node and allow it to deduplicate data on the client system

Register a node name of `JIM` and allow it to deduplicate data on the client system.

```
register node jim jimspass deduplication=clientorserver
```

Example: Register a node name of ED and set the role as a server-device for PVU estimation reporting

Register a node name of `ED` and set the role as a server-device for PVU estimation reporting.

```
register node ed pw459twx roleoverride=server
```

Example: Register a node on a source replication server

Define `NODE1` to a source replication server. Specify a replication rule for the backup data that belongs to `NODE1` so that active backup data is replicated with a high priority. Enable replication for the node.

```
register node node1 bkreplruledefault=active_data_high_priority replstate=enabled
```

Example: Register a node that authenticates with an LDAP server

Register a node name of `NODE17` that must authenticate with an LDAP server.


```
register node nodelpwd authentication=ldap
```

Tip: When you register a node in this way, an administrative user ID is not created.

Example: Register a node to communicate with a server by using strict session security

Register a node name of NODE4 to use the strictest security settings to authenticate with the server.

```
register node node4pwd sessionsecurity=strict
```

Example: Register a node and enable recovery of damaged files

Register a node name of PAYROLL. For the PAYROLL node, enable the recovery of damaged files from a target replication server.

```
register node payroll recoverdamaged=yes
```

Related commands

Table 2. Commands related to REGISTER NODE

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE MACHNODEASSOCIATION	Associates an IBM Spectrum Protect node with a machine.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
LOCK NODE	Prevents a client from accessing the server.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PVUESTIMATE	Displays an estimate of the client-devices and server-devices being managed.
QUERY REPLNODE	Displays information about the replication status of a client node.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REMOVE REPLNODE	Removes a node from replication.
RENAME NODE	Changes the name for a client node.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET DEDUPVERIFICATIONLEVEL	Specifies the percentage of extents verified by the server during client-side deduplication.

Command	Description
SET REPLRECOVERDAMAGED	Specifies whether node replication is enabled to recover damaged files from a target replication server.
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes that are associated with a client node.

Related concepts:

UNIX and Linux client root and authorized user tasks

Related reference:

Ssl client option

REMOVE commands

Use the REMOVE commands to remove an object from IBM Spectrum Protect™.

- REMOVE ADMIN (Delete an administrative user ID)
- **AIX** | **Linux** | **Windows** REMOVE DAMAGED (Remove damaged data from a source storage pool)
- REMOVE NODE (Delete a node or an associated machine node)
- REMOVE REPLNODE (Remove a client node from replication)
- REMOVE REPLSERVER (Remove a replication server)

REMOVE ADMIN (Delete an administrative user ID)

Use this command to remove an administrative user ID from the system.

You cannot remove the last system administrative user ID or the SERVER_CONSOLE administrative ID from the system.

For users of Lightweight Directory Access Protocol (LDAP) servers: The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REMOve Admin--admin_name--+-SYNCldapdelete-----No-----+-----<<
                               '-SYNCldapdelete-----+-No---+'
                               '-Yes-'
```

Parameters

admin_name (Required)

Specifies the administrative user ID to be removed.

SYNCldapdelete

Specifies whether to delete the administrative user ID on the Lightweight Directory Access Protocol (LDAP) server.

Yes

Deletes the administrative user ID on the LDAP server.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Does not delete the administrative user ID on the LDAP server. This is the default value.

Example: Remove an administrative user ID

Remove an administrative user ID larry that is not defined on an LDAP server. Issue the following command:

```
remove admin larry
```

Related commands

Table 1. Commands related to REMOVE ADMIN

Command	Description
LOCK ADMIN	Prevents an administrator from accessing IBM Spectrum Protect.
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
RENAME ADMIN	Changes an IBM Spectrum Protect administrator's name.

AIX Linux Windows

REMOVE DAMAGED (Remove damaged data from a source storage pool)

After storage pool conversion, use this command to remove damaged data from a storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL).

The REMOVE DAMAGED command permanently deletes damaged data from the storage pool.

Tip: Before you remove damaged data from the storage pool, try to recover an undamaged version of the data from a copy or active-data storage pool by issuing the RESTORE STGPOOL command. Recover an undamaged version of the data from a target replication server by issuing the REPLICATE NODE command and specifying the RECOVERDAMAGED=YES parameter.

Privilege class

To issue this command, you must have restricted storage privilege.

Syntax

```
>>-REMOve DAMAged--pool_name-- .-*----- .
| .-,----- . |
| v |
|---node_name--+'
.-Wait----No-----
>--+'----->>
'-Wait----+No--+'
'-Yes--'
```

Parameters

pool_name (Required)

Specify a primary storage pool that uses a FILE device class, a tape device class, or a virtual tape library (VTL). The storage pool contains the damaged data. This parameter is required.

node_name

Specifies the name of the client node. Separate multiple names with commas and no intervening spaces. You can use a wildcard character instead of a node name if you want to remove damage from all of the nodes in the storage pool.

Wait

Specifies whether to wait for the server to remove damaged data from the storage pool. This parameter is optional. The default value is NO. You can specify this parameter only from an administrative command line. You can specify one of the

following values:

No

Specifies that the command processes run in the background.

Yes

Specifies that the command processes run in the foreground. Messages are not displayed until the command completes processing.

Example: Remove damaged data from a storage pool and wait for the server to complete processing

Remove damaged data from a storage pool that is named POOL1 and wait for the server to complete processing in the foreground.

```
remove damaged pool1 wait=yes
```

Table 1. Commands related to REMOVE DAMAGED

Command	Description
CONVERT STGPOOL	Convert a storage pool to a directory-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
REPAIR STGPOOL	Repairs a directory-container storage pool.

REMOVE NODE (Delete a node or an associated machine node)

Use this command to remove a node from the server. If you are using disaster recovery manager and the node to be removed is associated with a machine, the association between the node and the machine is also deleted.

If a node is part of a collocation group and you remove the node from the server, the node is removed from the collocation group. If a node is removed and the node contained file spaces in a file space collocation group, those file spaces are removed from the group member list.

If you remove a node that stored data in a deduplicated storage pool, the node name DELETED is displayed in the QUERY OCCUPANCY command output until all data deduplication dependencies are removed.

When a node is removed, the corresponding administrative ID is removed only if the following issues are true:

- The administrator name is identical to the node name.
- The administrator has client owner or client access authority *only* to the node that is being removed.
- The administrator is not a managed object.

Before you can remove a node, you must delete all backup and archive file spaces that belong to that node.

Before you can remove a NAS node that has a corresponding data mover, you must complete the following tasks in order:

1. Delete any paths from the data mover
2. Delete the data mover
3. Delete all virtual file space definitions for the node
4. Remove the NAS node

For users of Lightweight Directory Access Protocol (LDAP) servers: The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-REMove Node--node_name--+-SYNCldapdelete-----No----->>  
'-SYNCldapdelete-----+-No-----'
```

Parameters

node_name (Required)

Specifies the name of the node to be removed.

SYNCDapdelete

Specifies whether to remove the node from the Lightweight Directory Access Protocol (LDAP) server.

Yes

Specifies that the node is removed.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Specifies that the node is not removed. This is the default value.

Example: Remove a client node

Remove the client node LARRY.

```
remove node larry
```

Related commands

Table 1. Commands related to REMOVE NODE

Command	Description
AIX Windows DELETE MACHNODEASSOCIATION	AIX Windows Deletes association between a machine and node.
DELETE DATAMOVER	Deletes a data mover.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
DELETE PATH	Deletes a path from a source to a destination.
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
LOCK NODE	Prevents a client from accessing the server.
QUERY COLLOGGROUP	Displays information about collocation groups.
AIX Windows QUERY MACHINE	AIX Windows Displays information about machines.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY SESSION	Displays information about all active administrator and client sessions with IBM Spectrum Protect.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RENAME NODE	Changes the name for a client node.

REMOVE REPLNODE (Remove a client node from replication)

Use this command to remove a node from replication if you no longer want to replicate the data that belongs to the node.

You cannot delete client node data by issuing the REMOVE REPLNODE command. You can issue the command on a source or on a target replication server. You can only issue this command from an administrative command-line client. You cannot issue this command from the server console.

If you issue the REMOVE REPLNODE command for a client node whose replication mode is set to SEND or RECEIVE, the mode is set to NONE. The replication state is also set to NONE. After you remove a client node from replication, the target replication

server can accept backup, archive, and space-managed data directly from the node.

If a client node is removed from replication, information in the database about replication for the node is deleted. If the client node is enabled for replication later, the replication process replicates all the data that is specified by replication rules and settings.

When you issue the REMOVE REPLNODE command, the data that belongs to a client node is not deleted. To delete file space data that belongs to the client node, issue the DELETE FILESPACE command for each of the file spaces that belong to the node. If you do not want to keep the client node definition, issue the REMOVE NODE command. To delete file space data and the client node definition, issue DELETE FILESPACE and REMOVE NODE on the target replication server.

Restriction: If a node replication process is running for a client node that is specified by this command, the command fails and the replication information for the node is not removed.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
      .-|-----|
      v      |
>>-REMOve REPLNode-----+--node_name-----+--+-----><
                          '-node_group_name-'
```

Parameters

node_name or node_group_name (Required)

Specifies the name of the client node or defined group of client nodes that you want to remove from replication. To specify multiple client node names and client-node group names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify client node names, but not to specify client-node group names. You cannot combine node or node group names with the domain name.

Example: Remove three client nodes and a client node group from replication

The names of the client nodes are NODE1, NODE2, and NODE3. The name of the client node group is PAYROLL. Issue the following command on the source and target replication servers:

```
remove replnode node*,payroll
```

Related commands

Table 1. Commands related to REMOVE REPLNODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.

REMOVE REPLSERVER (Remove a replication server)

Use this command to remove or to switch to a replication server from the list of replication servers. This command deletes all information about replication state for all nodes that were replicated to that server.

You can issue the command on a source or on a target replication server.

Restriction: You cannot delete client node data by using the REMOVE REPLSERVER command.

Use the command to switch replication servers and to remove replication information for an old server. The command does not affect the current replication mode or state of any node definitions. Issue the command on both the source and target servers to keep the replication state information about both servers consistent.

Restriction: If you specify the default replication server for the REMOVE REPLSERVER command and a node replication process is running, the command fails and no replication information is removed.

This command runs as a background operation and it cannot be canceled. IBM Spectrum Protect™ deletes replication information that is associated with the specified server as a series of batch database transactions. If a system failure occurs, a partial deletion can occur.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REMOve REPLServer--GUID-----<<
```

Parameters

replication_guid (Required)

The unique identifier for the replication server that is being removed. You can use wildcards to specify the Replication Global Unique Identifier (GUID), however, only one GUID can match the wildcard. If the wildcard sequence matches more than one GUID, the command fails. You must qualify the wildcard string until only the GUID that you want to delete is found.

Example: Use a wildcard to remove a replication server

Remove a replication server by using a wildcard character to indicate the GUID.

```
remove replserver e*
```

Related commands

Table 1. Commands related to REMOVE REPLSERVER

Command	Description
REMOVE REPLNODE (Remove a client node from replication)	Removes a node from replication.
QUERY REPLSERVER (Query a replication server)	Displays information about replicating servers.

RENAME commands

Use the RENAME commands to change the name of an existing object.

- RENAME ADMIN (Rename an administrator)
- RENAME FILESPACE (Rename a client file space on the server)
- RENAME NODE (Rename a node)
- RENAME SCRIPT (Rename an IBM Spectrum Protect script)
- RENAME SERVERGROUP (Rename a server group)
- RENAME STGPOOL (Change the name of a storage pool)

RENAME ADMIN (Rename an administrator)

Use this command to change an administrative user ID. Existing information for this administrator such as password, contact information, and privilege classes is not altered.

If you assign an existing administrative user ID to another person, use the UPDATE ADMIN command to change the password.

When an administrator and a node share a name and you change the administrator authentication method, the node authentication method also changes. If you rename an administrator to the same name as an existing node, the authentication

method and the SSLREQUIRED setting for the node can change. If those settings are different, after the renaming, both administrator and node will have the same authentication method and SSLREQUIRED setting.

For users of Lightweight Directory Access Protocol (LDAP) servers:

- The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.
- Do not rename an administrative user ID to match a node name. If the names match, you might see unexpected behavior because of automatic password changes that update the same password twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update might fail.

You cannot rename the SERVER_CONSOLE administrative ID.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REName Admin--current_admin_name--new_admin_name----->
. -SYNClapdelete---No-----
>--+-----+-----><
' -SYNClapdelete---+No--+-'
      '-Yes-'
```

Parameters

current_admin_name (Required)

Specifies the administrative user ID to be renamed.

new_admin_name (Required)

Specifies the new administrative user ID. The maximum length of the name is 64 characters.

SYNClapdelete

Specifies whether to delete the administrative user ID on the Lightweight Directory Access Protocol (LDAP) server and replace the ID with a new one.

Yes

Deletes the administrative user ID on the LDAP server and replaces it with a new ID.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Does not delete and replace the administrative user ID on the LDAP server. This is the default value.

Example: Rename an administrator

Rename the IBM Spectrum Protect administrator CLAUDIA to BILL.

```
rename admin claudia bill
```

Related commands

Table 1. Commands related to RENAME ADMIN

Command	Description
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.

RENAME FILESPACE (Rename a client file space on the server)

Use this command to rename an existing client file space on the server to a new file space name or to rename imported file spaces.

You might want to rename a file space that was imported or to cause the creation of new Unicode-enabled file spaces for Unicode-enabled clients.

Restriction: Do not rename NAS or VMware file spaces. If you rename a NAS or VMware file space, it is no longer visible and cannot be restored. To restore a renamed NAS or VMware file space, you must rename it back to its original name and set the force parameter as follows:force=yes

Privilege class

Any administrator with unrestricted policy authority or with restricted policy authority over the client's policy domain can issue this command.

Syntax

```
>>-REName Filespace--node_name----->
>--current_file_space_name--new_file_space_name----->
.-NAMEType---SERVER-----
>+-----+-----+----->
'-NAMEType---+--SERVER--+'
      +-UNICODE-+
      '-FSID----'

.-NEwNAMEType---SERVER-----
>+-----+-----+-----><
|                                     (1) | '-force---yes-'
'-NEwNAMEType---+--UNICODE-----+'
      '-HEXadecimal-'
```

Notes:

1. This parameter is the default when you specify NAMEType=UNICODE.

Parameters

node_name (Required)

Specifies the name of the client node to which the file space to be renamed belongs.

current_file_space_name (Required)

Specifies the name of the file space to be renamed. A file space name is case-sensitive and must be specified exactly as defined to the server. Virtual file space mapping names are allowed.

new_file_space_name (Required)

Specifies the new name for the file space. A client file space name is case-sensitive and must be specified exactly as it is to be defined to the server. This parameter cannot be an existing virtual file space mapping name. If the current_file_space_name is a virtual file space, the new_file_space_name must follow all the rules for defining a virtual file space name. See the DEFINE VIRTUALFSMAPPING command for more information.

Important: If the new name type is hexadecimal, specify valid UTF-8 hexadecimal values so the server's code page displays the file space name as intended. For example, do not specify a value that can be interpreted as a backspace. When you rename a file space that is part of a file space collocation group, the collocation group is updated with the new name.

NAMEType

Specify how you want the server to interpret the current file space name that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect™ clients with Windows, Macintosh OS X, and NetWare operating systems.

The default value is SERVER. If a virtual file space mapping name is specified, you must use SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space name.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space name as the file space ID (FSID).

NEWNAMETYPE

Specify how you want the server to interpret the new file space name that you enter. The default is SERVER if you specified the NAMETYPE as SERVER, or if the file space to be renamed is not Unicode. The default is UNICODE if you specified the NAMETYPE as UNICODE, or if the file space to be renamed is Unicode. If a virtual file space mapping name is specified, you must use SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space name.

UNICODE

The server converts the file space name that is entered from the server code page, to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. If the conversion is not successful, you might want to specify the HEXADECIMAL parameter.

HEXadecimal

The server interprets the file space name that you enter as the hexadecimal representation of a name in Unicode. Using hexadecimal ensures that the server is able to correctly rename the file space, regardless of the server's code page.

To view the hexadecimal representation of a file space name, you can use the QUERY FILESPACE command with FORMAT=DETAILED.

Restriction: You cannot specify a new name of a type that is different from the original name. You can rename a file space that is Unicode to another name in Unicode. You can rename a file space that is not Unicode, and use a new name in the server's code page. You cannot mix the two types.

force

To rename a NAS or VMware file space you must set this parameter as follows: force=yes

Rename an imported file space to prevent overwriting

An AIX® client node named LARRY backed up file space /r033 to the IBM Spectrum Protect server. The file space was exported to tape and later reimported to the server. When this file space was imported, a system-generated name, /r031, was created for it because /r033 existed for client node LARRY.

Client node LARRY, however, already had a file space named /r031 that was not backed up, therefore, was unknown to the server. Unless the imported file space is renamed, it overlays file space /r031 because the file space name generated by the IMPORT function is the same as a file space on client node LARRY that is unknown to the server.

Use the following command to rename imported file space /r031. The new name, /imported-r033, identifies that the new file space is an imported image of file space /r033.

```
rename filesystem larry /r031 /imported-r033
```

Rename file space to create a Unicode-enabled file space

Client JOE is using an English Unicode-enabled IBM Spectrum Protect client. JOE backed up several large file spaces that are not Unicode that is enabled in server storage. File space \\joe\c\$ contains some files with Japanese file names that cannot be backed up to a file space that is not Unicode that is enabled. Because the file spaces are large, the administrator does not want to convert all of JOE's file spaces to Unicode-enabled file spaces now. The administrator wants to rename only the non-Unicode file space, \\joe\c\$, so that the next backup of the file space causes the creation of a new Unicode-enabled file space. The new Unicode-enabled file space allows the Japanese files to be successfully backed up.

Use the following command to rename \\joe\c\$:

```
rename filesystem joe \\joe\c$ \\joe\c$_old
```

Related commands

Table 1. Commands related to RENAME FILESPACE

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
EXPORT NODE	Copies client node information to external media or directly to another server.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY OCCUPANCY	Displays file space information by storage pool.

RENAME NODE (Rename a node)

Use this command to rename a node.

If you are assigning an existing node ID to another person, use the UPDATE NODE command to change the password.

For users of Lightweight Directory Access Protocol (LDAP) servers:

- The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.
- Do not rename a node to match an existing administrative user ID. If you rename a node, and the node name matches an administrative user ID, you might see unexpected behavior because of automatic password changes that update the same password twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update might fail.

Restrictions:

- You cannot rename a NAS node name that has a corresponding data mover defined. If the data mover has defined paths, the paths must first be deleted.
- If a node is configured for replication, it cannot be renamed.

If you rename a node to the same name as an existing administrator, the administrator authentication method and SSLREQUIRED setting are updated to match the node. When a node and an administrator share a name and you change the node authentication method or the node SSLREQUIRED setting, the administrator settings also change. You must have system level authority to update the node authentication method or the node SSLREQUIRED setting and also update a same-named administrator.

Privilege class

You must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-REName Node--current_node_name--new_node_name----->
      .-SYNCldapdelete---No-----
>-----+-----+-----><
      '-SYNCldapdelete---+No---+'
              '-Yes-'
```

Parameters

current_node_name (Required)

Specifies the name of the node to be renamed.

new_node_name (Required)

Specifies the new name of the node. The maximum length is 64 characters.

SYNCDapdelete

Specifies whether the node name is deleted and replaced on the Lightweight Directory Access Protocol (LDAP) server.

Yes

Specifies that the node name is deleted and replaced.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Specifies that the node name is not deleted and replaced. This is the default value.

Example: Rename a node

Rename the node JOE to JOYCE.

```
rename node joe joyce
```

Example: Rename a node that shares a namespace with other servers

Rename the node JOYCE to JOE and do not delete the previous name from corresponding LDAP servers.

```
rename node joyce joe
```

Related commands

Table 1. Commands related to RENAME NODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
UPDATE NODE	Changes the attributes that are associated with a client node.

Related tasks:

Managing NAS file server nodes

RENAME SCRIPT (Rename an IBM Spectrum Protect script)

Use this command to rename an IBM Spectrum Protect™ script.

Privilege class

To issue this command, you must have operator, policy, system, storage, or system privilege.

Syntax

```
>>-REName SCRIPT--current_script_name--new_script_name -----><
```

Parameters

current_script_name (Required)

Specifies the name of the script to rename.

new_script_name (Required)

Specifies the new name for the script. The name can contain as many as 30 characters.

Example: Rename a script

Rename SCRIPT1 to a new script named SCRIPT2.

```
rename script script1 script2
```

Related commands

Table 1. Commands related to RENAME SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

RENAME SERVERGROUP (Rename a server group)

Use this command to rename a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REName SERVERGroup--current_group_name--new_group_name-----<<
```

Parameters

current_group_name (Required)

Specifies the server group to rename.

new_group_name (Required)

Specifies the new name of the server group. The maximum length of the name is 64 characters.

Example: Rename a server group

Rename server group WEST_COMPLEX to BIG_WEST.

```
rename servergroup west_complex big_west
```

Related commands

Table 1. Commands related to RENAME SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVERGROUP	Displays information about server groups.
UPDATE SERVERGROUP	Updates a server group.

RENAME STGPOOL (Change the name of a storage pool)

Use this command to change the name of a storage pool. You can change storage pool names to use the same names on a configuration manager and its managed servers.

When you rename a storage pool, any administrators with restricted storage privilege for the old storage pool automatically retain restricted storage privilege for the renamed storage pool. If the renamed storage pool is in a storage pool hierarchy, the hierarchy

is preserved. You must update the management class or copy group to specify the new storage pool name as the destination for files.

If processes are active when a storage pool is renamed, the old name might still be displayed in messages or queries for those processes.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-REName STGpool--current_pool_name--new_pool_name-----<<
```

Parameters

current_pool_name (Required)

Specifies the storage pool to rename.

new_pool_name (Required)

Specifies the new name of the storage pool. The maximum length of the name is 30 characters.

Example: Change the name of a storage pool

Rename storage pool STGPOOLA to STGPOOLB:

```
rename stgpool stgpoola stgpoolb
```

Related commands

Table 1. Commands related to RENAME STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE STGPOOL	Deletes a storage pool from server storage.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.

AIX Linux Windows

REPAIR STGPOOL (Repair a directory-container storage pool)

Use this command to repair deduplicated extents in a directory-container storage pool. Damaged deduplicated extents are repaired with extents that are backed up to the target replication server or to container-copy storage pools on the same server.

Restrictions:

- You can issue the REPAIR STGPOOL command only if you already issued the PROTECT STGPOOL command to back up data to another storage pool on a replication target server or on the same server.
- When you repair a directory-container storage pool from the replication server, the REPAIR STGPOOL command fails when any of the following conditions occur:
 - The target server is unavailable.
 - The target storage pool is damaged.
 - A network outage occurs.
- When you repair a directory-container storage pool from container-copy pools, the REPAIR STGPOOL command fails when any of the following conditions occur:

- o The container-copy storage pool is unavailable.
- o The container-copy storage pool is damaged.

Privilege class

To issue this command, you must have system privilege.

Syntax when the source is the replication server

```

                .-SRCLOCation---Replserver-.
>>-REPAir STGPool--pool_name-----+----->
                '-SRCLOCation---Replserver-'

                .-MAXSESSions---1-----
>--+-----+----->
                '-MAXSESSions-----number_sessions--'

                .-Preview---No----- .-Wait---No-----
>--+-----+----->>
                '-Preview---+No--+-' '-Wait---+No--+-'
                    '-Yes-'           '-Yes-'

```

Syntax when the source is a storage pool on the same server

```

>>-REPAir STGPool--pool_name--SRCLOCation---Local----->

                .-Preview---No----- .-Wait---No-----
>--+-----+----->>
                '-Preview---+No--+-' '-Wait---+No--+-'
                    '-Yes-'           '-Yes-'

```

Parameters

pool_name (Required)

Specifies the name of the directory-container storage pool that contains the data that must be repaired.

SRCLOCation

Specifies the source location that is used to repair the data. The default value is REPLSERVER. This parameter is only required when the source location is on the same server. You can specify one of the following values:

Local

Specifies that the data is repaired from container-copy storage pools on the same server.

Replserver

Specifies that the data is repaired from a directory-container storage pool on the target replication server.

MAXSESSions

Specifies the maximum number of data sessions that can send data to a target server. This parameter is optional when you repair data from a replication server.

The value that you specify can be in the range 1 - 20. The default value is 1. If you increase the number of sessions, you can repair the storage pool faster.

When you set a value for the MAXSESSIONS parameter, ensure that the available bandwidth and the processor capacity of the source and target servers are sufficient.

Tips:

- If you issue a QUERY SESSION command, the total number of sessions might exceed the number of data sessions.
- The number of sessions that are used to repair storage pools depends on the amount of data that is repaired. If you repair only a small amount of data, there is no benefit to increasing the number of sessions.

Preview

Specifies whether to preview data or to repair the data. This parameter is optional. The default value is NO. You can specify one of the following values:

No
Specifies that the data is repaired to the storage pool but the data is not previewed.

Yes
Specifies that the data is previewed but not repaired.

Wait

Specifies whether to wait for the server to complete the repair processing of the storage pool. This parameter is optional. The default value is NO. You can specify this parameter only from an administrative command line. You can specify one of the following values:

No
Specifies that the command processes run in the background. To monitor the background processing of the REPAIR STGPOOL command, issue the QUERY PROCESS command.

Yes
Specifies that the command processes run in the foreground. Messages are not displayed until the command completes processing.

Example: Repair a storage pool and preview the data

Repair a storage pool that is named POOL1 and preview the data.

```
repair stgpool pool1 preview=yes
```

Example: Repair a storage pool and specify a maximum number of sessions

Repair a storage pool that is named POOL1 and specify 10 maximum sessions.

```
repair stgpool pool1 maxsessions=10
```

Example: Repair a storage pool from tape

Repair a storage pool that is named POOL1 and specify local for the source location.

```
repair stgpool pool1 SRCLOCation=local
```

Table 1. Commands related to REPAIR STGPOOL

Command	Description
DEFINE STGPOOL (directory-container)	Define a directory-container storage pool.
DEFINE STGPOOL (container-copy)	Define a container-copy storage pool that stores copies of data from a directory-container storage pool.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.

REPLICATE NODE (Replicate data in file spaces that belong to a client node)

Use this command to replicate data in file spaces that belong to one or more client nodes or defined groups of client nodes.

When you issue this command, a process is started in which data that belongs to the specified client nodes is replicated according to replication rules. Files that are no longer stored on the source replication server, but that exist on the target replication server, are deleted during this process.

Tip: Avoid conflicts in managing administrative IDs and client option sets by identifying the IDs and option sets that are replicated to the target server and the IDs and option sets that are managed in an enterprise configuration. You cannot define an administrative user ID for a registered node if an administrative ID exists for the same node.

If a node replication process is already running for a client node that is specified by this command, the node is skipped, and replication begins for other nodes that are enabled for replication.

After the node replication process is completed, a recovery process can be started on the target replication server. Files are recovered only if all the following conditions are met:

- Version 7.1.1 or later, is installed on the source and target replication servers.
- The REPLRECOVERDAMAGED system parameter is set to ON. The system parameter can be set by using the SET REPLRECOVERDAMAGED command.
- The source server includes at least one file that is marked as damaged in the node that is being replicated.
- The node data was replicated before the damage occurred.

The following table describes how settings affect the recovery of damaged, replicated files.

Restriction: You cannot use the REPLRECOVERDAMAGED parameter for directory-container or cloud storage pools.

Table 1. Settings that affect the recovery of damaged files

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result
OFF	YES, NO, or not specified	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
OFF	ONLY	YES or NO	An error message is displayed because files cannot be recovered when the REPLRECOVERDAMAGED system parameter is set to OFF.
ON	YES	YES or NO	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	NO	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
ON	ONLY	YES or NO	Damaged files are recovered from the target replication server, but standard node replication does not occur.
ON	Not specified	YES	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	Not specified	NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.

Tip: When the QUERY PROCESS command is issued during node replication, the output can show unexpected results for the number of completed replications. The reason is that, for node replication purposes, each file space is considered to contain three logical file spaces:

- One for backup objects
- One for archive objects
- One for space-managed objects

By default, the QUERY PROCESS command generates results for each logical file space. Other factors also affect the output of the QUERY PROCESS command:

- If a file space has a replication rule that is set to NONE, the file space is not included in the count of file spaces that are being processed.
- If you specify data types in the REPLICATE NODE command, only those data types are included in the count of file spaces that are being processed, minus any file spaces that are excluded.


```

'-FORCEREconcile-----+No-----'
                        +-Yes--+
                        '-FULL-'

.-TRANSFERMethod-----Tcip-----
>-----+-----+-----><
'-TRANSFERMethod-----+Tcip-----'
                        |         (3) |
                        '-Fasp-----'

```

Notes:

1. Do not mix file space identifiers (FSIDs) and file space names in the same command.
2. Do not specify FSID if you use wildcard characters for the client node name.
3. **Linux** The TRANSFERMETHOD parameter is available only on Linux x86_64 operating systems.

Parameters

node_name or node_group_name (Required)

Specifies the name of the client node or defined group of client nodes whose data is to be replicated. You can also specify a combination of client node names and client-node group names. To specify multiple client node names or client-node group names, separate the names with commas with no intervening spaces. You can use wildcard characters with client node names but not with client-node group names. The replication rules for all file spaces in the specified client nodes are checked.

filesystem_name or FSID

Specifies the name of the file space or the file space identifier (FSID) to be replicated. A name or FSID is optional. If you do not specify a name or an FSID, all the data in all the file spaces for the specified client nodes is eligible for replication.

filesystem_name

Specifies the name of the file space that has data to be replicated. File space names are case-sensitive. To determine the correct capitalization for the file space, issue the QUERY FILESPACE command. Separate multiple names with commas with no intervening spaces. When you specify a name, you can use wildcard characters.

A server that has clients with file spaces that are enabled for Unicode might have to convert the file space name. For example, the server might have to convert a name from the server code page to Unicode. For details, see the NAMETYPE parameter. If you do not specify a file space name, or if you specify a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

FSID

Specifies the file space identifier for the file space to be replicated. The server uses FSIDs to find the file spaces to replicate. To determine the FSID for a file space, issue the QUERY FILESPACE command. Separate multiple FSIDs with commas with no intervening spaces. If you specify an FSID, the value of the NAMETYPE parameter must be FSID.

NAMETYPE

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for IBM Spectrum Protect™ clients that are enabled for Unicode and that have Windows, Macintosh OS X, or NetWare operating systems.

Use this parameter only when you enter a partly qualified or fully qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret file space names.

UNICODE

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines.

FSID

The server interprets file space names by using their file space identifiers.

CODETYPE

Specifies the type of file spaces to be included in node replication processing. Use this parameter only when you enter a single wildcard character for the file space name. The default value is BOTH, which specifies that file spaces are included regardless of code page type. You can specify one of the following values:

UNICODE

Specifies file spaces that are only in Unicode.

NONUNICODE

Specifies file spaces that are not in Unicode.

BOTH

Specifies all file spaces regardless of code page type.

DATATYPE

Specifies the type of data to be replicated. Data is replicated according to the replication rule that applies to the data type. This parameter is optional. You can specify one or more data types. If you do not specify a data type, all backup, archive, and space-managed data is replicated. Separate multiple data types with commas with no intervening spaces. You cannot use wildcard characters. You can specify one of the following values:

ALL

Replicates all backup, archive, and space-managed data in a file space according to the rule that is assigned to the data type. For example, suppose that NODE1 has a single file space. The following replication rules apply:

- The file space rules for backup and archive data in the file space are set to ALL_DATA.
- The file space rule for space-managed data is set to DEFAULT.
- The client node rule for space-managed data is set to NONE.

If you issue `REPLICATE NODE NODE1 DATATYPE=ALL`, only backup data and archive data are replicated.

BACKUP

Replicates active and inactive backup data in a file space if the controlling replication rule is ALL_DATA, ACTIVE_DATA, ALL_DATA_HIGH_PRIORITY, or ACTIVE_DATA_HIGH_PRIORITY.

BACKUPActive

Replicates only active backup data in a file space if the controlling replication rule is ACTIVE_DATA or ACTIVE_DATA_HIGH_PRIORITY.

ARCHIVE

Replicates archive data only in a file space if the controlling replication rule is ALL_DATA or ALL_DATA_HIGH_PRIORITY.

SPACEManaged

Replicates only space-managed data in a file space if the controlling replication rule is ALL_DATA or ALL_DATA_HIGH_PRIORITY.

PRIORITY

Specifies the data to replicate based on the priority of the replication rule. You can specify one of the following values:

All

Replicates all data in a file space if the controlling replication rule is ALL_DATA, ACTIVE_DATA, ALL_DATA_HIGH_PRIORITY, or ACTIVE_DATA_HIGH_PRIORITY.

High

Replicates only data in a file space that has a controlling replication rule of ALL_DATA_HIGH_PRIORITY or ACTIVE_DATA_HIGH_PRIORITY.

Normal

Replicates only data in a file space that has a controlling replication rule of ALL_DATA or ACTIVE_DATA.

MAXSESSIONS

Specifies the maximum allowable number of data sessions to use for sending data to a target replication server. This parameter is optional. The value can be 1 - 99. The default value is 10.

Increasing the number of sessions can improve node replication throughput.

When you set this value, consider the number of logical and physical drives that can be dedicated to the replication process. To access a sequential-access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on the following factors:

- Other IBM Spectrum Protect and system activity
- The mount limits of the device classes for the sequential access storage pools that are involved

Ensure that sufficient mount points and drives are available to allow node replication processes to complete. Each replication session might need a mount point on the source and target replication servers for storage pool volumes. If the device type is not FILE, each session might also need a drive on both the source and target replication servers.

When you set a value for MAXSESSIONS, also consider the available bandwidth and the processor capacity of the source and target replication servers.

Tip:

- The value that is specified by the MAXSESSIONS parameter applies only to data sessions. Data sessions are sessions during which data is sent to a target replication server. However, if you issue a QUERY SESSION command, the total number of sessions might exceed the number of data sessions. The difference is because of short control sessions that are used for querying and setting up replication operations.
- The value of the MAXSESSIONS parameter represents the maximum allowable number of sessions. The number of sessions that are used for replication depends on the amount of data to be replicated. If you are replicating only a small amount of data, you do not achieve any benefit by increasing the number of sessions. The total number of sessions might be less than the value that is specified by the MAXSESSIONS parameter.

Preview

Specifies whether to preview data. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the data is replicated to the target server but that the data is not previewed.

Yes

Specifies that data is previewed but not replicated. If you specify PREVIEW=YES, only volumes that must be physically mounted, such as tape volumes, are displayed. Volumes that are assigned to storage pools that have a device class of FILE are not displayed.

The following information is displayed in the output:

- The names of client nodes whose data would be replicated.
- The number of files that would be replicated or deleted.
- The estimated amount of time it would take to complete the node replication process.
- A list of volumes that would be mounted.
- A summary of information about replicated, damaged data. The summary lists the number of nodes, file spaces, files, and bytes that can be recovered during a replication recovery process. The summary is displayed only if RECOVERDAMAGED=YES or RECOVERDAMAGED=ONLY is specified.

If the client node data that is specified by the REPLICATE NODE command was never replicated and you specify PREVIEW=YES, the node and its file spaces are automatically defined on the target replication server.

LISTfiles

Specifies whether to list the names of files that would be replicated. This parameter is optional. The default is NO. Specifying this parameter signifies that the WAIT parameter is set to YES and that you cannot issue the WAIT parameter from the server console.

You can specify one of the following values:

No

Specifies that the names of files that would be replicated are not displayed.

Yes

Specifies that the names of files that would be replicated are displayed.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the command processes in the background. To monitor the background processing of the REPLICATE NODE command, issue the QUERY PROCESS command.

Yes

Specifies that the command processes in the foreground. Messages are not displayed until the command completes processing. You cannot specify WAIT=YES from the server console.

RECOVERDamaged

Specifies whether a recovery process is started on a target replication server after the node replication process is completed. This parameter is optional, and it overrides any value that you specified for the RECOVERDamaged parameter when you defined or updated a node. You can specify one of the following values:

Yes

Specifies that a replication process is started to recover damaged files, but only if the setting for the REPLRECOVERDAMAGED system parameter is ON. If the setting is OFF, damaged files are not recovered.

No

Specifies that damaged files are not recovered.

Only

Specifies that a replication process is started for the sole purpose of recovering damaged files, but only if the setting for the REPLRECOVERDAMAGED system parameter is ON. If the setting is OFF, damaged files are not recovered, and you receive a notification that recovery was not started.

Restriction: If you specify an invalid combination of values and settings for file recovery, replication is stopped, and an error message is displayed.

FORCEREconcile

Specifies whether to compare all files on the source replication server with files on the target replication server and to synchronize the differences between them. Before V7.1.1, this behavior was the default for replication processing. When IBM® Tivoli® Storage Manager V7.1.1 or later is installed on the source and target replication servers, a reconcile is automatically completed during initial replication. After initial replication, you might use this parameter for the following reasons:

- To synchronize files on the source and target replication servers if they are different.
- To replicate inactive files that were skipped after you change your replication rules from ACTIVE_DATA to ALL_DATA.
- To delete inactive files from the target replication server when you change your replication rules from ALL_DATA to ACTIVE_DATA.
- To ensure that you replicate only active data when you are using the ACTIVE_DATA replication rule so that the target replication server has active files only.
- To resynchronize the files so that the target replication server has the same files as the source replication server if you have previously or are currently using the policies on the target replication server to manage replicated files.
- To resynchronize the files on the source and target replication servers if the database is regressed to an earlier point-in-time by using a method other than the DSMSEV RESTORE DB command.
- To rebind files to the new management class on the target replication server if this management class did not exist when the files were replicated. You must be using the policies that are defined on the target replication server to manage replicated files.
- To remove all files on a target server for a node and file space that do not exist on the replication source server.

Remember: When the ACTIVE_DATA rule is assigned, a reconcile is completed only for active files on the source replication server.

This parameter is optional. You can specify one of the following values:

No

Specifies that replication processing does not force a reconcile to compare all files on the source replication server with files on the target replication server. Instead, replication processing tracks file changes on the source replication server since the last replication and synchronizes these changes on the target replication server. NO is the default value.

Yes

Specifies that replication processing forces a reconcile to compare all files on the source replication server with files on the target replication server and synchronizes the files on the target replication server with the source replication server.

FULL

Specifies that replication processing forces a reconcile to compare all files on the source replication server with files on the target replication server and synchronizes the files on the target replication server with the source replication server. Any files that do not exist on the source replication server are removed from the target replication server.

Files might be removed for the following reasons:

- As a result of file space backup or import operations, files on the target replication server are no longer managed by replication processing.
- Replication-related orphaned objects on the target server are no longer managed by replication processing.

Restriction: Objects are deleted from the target replication server when nodes and file spaces are recognized by a replication process but the objects are not recognized.

Linux TRANSFERMethod

Linux Specifies the method that is used for server-to-server data transfer. This parameter is optional. You can specify one of the following values:

Tcpip

Specifies that TCP/IP is used to transfer data. This value is the default.

Fasp

Specifies that Aspera® Fast Adaptive Secure Protocol (FASP) technology is used to transfer data. Aspera FASP technology can help you optimize data transfer in a wide area network (WAN). If you specify TRANSFERMETHOD=FASP, you override any TRANSFERMETHOD parameters that you specified on the DEFINE SERVER or UPDATE SERVER commands.

Restrictions:

- Only data that is stored in a directory-container storage pool can be transferred by using Aspera FASP technology. Data that is not stored in a directory-container storage pool is transferred by using TCP/IP.
- Before you enable Aspera FASP technology, determine whether the technology is appropriate for your system environment and install the appropriate licenses. For instructions, see Determining whether Aspera FASP technology can optimize data transfer in your system environment. If the licenses are missing or expired, node replication fails.
- If WAN performance meets your business needs, do not enable Aspera FASP technology.

Example: Replicate data by data type and priority

Replicate high-priority active backup data and high-priority archive data that belongs to all the client nodes in group PAYROLL.

```
replicate node payroll datatype=backupactive,archive priority=high
```

Example: Replicate all the data that belongs to a node according to the assigned replication rules

NODE1 has a single file space. The following replication rules apply:

- File space rules:
 - Backup data: ACTIVE_DATA
 - Archive data: DEFAULT
 - Space-managed data: DEFAULT
- Client node rules:
 - Backup data: DEFAULT
 - Archive data: ALL_DATA_HIGH_PRIORITY
 - Space-managed data: DEFAULT
- Server rules:
 - Backup data: ALL_DATA
 - Archive data: ALL_DATA
 - Space-managed data: NONE

```
replicate node node1 priority=all
```

Active backup data in the file space is replicated with normal priority. Archive data is replicated with high priority. Space-managed data is not replicated.

Example: Recover damaged files without starting the full replication process

Without starting the full replication process, recover any damaged files in the client nodes of the PAYROLL group. Ensure that the setting for the REPLRECOVERDAMAGED system parameter is ON. Then, issue the following command:

```
replicate node payroll recoverdamaged=only
```

Related commands

Table 2. Commands related to REPLICATE NODE

Command	Description
CANCEL PROCESS	Cancels a background server process.
CANCEL REPLICATION	Cancels node replication processes.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY REPLRULE	Displays information about node replication rules.
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE REPLNODE	Removes a node from replication.
AIX Linux Windows PROTECT STGPOOL	Protects a directory-container storage pool.
SET REPLRECOVERDAMAGED	Specifies whether node replication is enabled to recover damaged files from a target replication server.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes that are associated with a client node.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

REPLY (Allow a request to continue processing)

Use this command and an identification number to inform the server that you have completed a requested operation. Not all server requests require a reply. This command is required only if the request message specifically indicates that a reply is needed.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
>>-REPLY--request_number--+-+-----+-----<<
                        '-LABEL-----volume_label-'
```

Parameters

request_number (Required)

Specifies the identification number of the request.

LABEL

Specifies the label to be written on a volume when you reply to a message from a LABEL LIBVOLUME command process. This parameter is optional.

Example: Reply to a request

Respond to a reply request using 3 as the request number.

Related commands

Table 1. Commands related to REPLY

Command	Description
CANCEL REQUEST	Cancels pending volume mount requests.
QUERY REQUEST	Displays information about all pending mount requests.

RESET PASSEXP (Reset password expiration)

Use the RESET PASSEXP command to reset the password expiration period to the common expiration period for administrator and client node passwords. The RESET PASSEXP command does not apply to passwords that are stored on an LDAP directory server.

Restriction: You cannot reset the password expiration period to the common expiration period with the SET PASSEXP command.

Use the QUERY STATUS command to display the common password expiration period.

Restriction: If you do not specify either the NODE or ADMIN parameters, the password expiration period for all client nodes and administrators will be reset.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-RESet PASSExp----->
      |               .-,----- . |
      |               v             | |
      |'-Node-----node_name-+-'|
>-----><
      |               .-,----- . |
      |               v             | |
      |'-Admin-----admin_name-+-'|
```

Parameters

Node

Specifies the name of the node whose password expiration period you would like to reset. To specify a list of nodes, separate the names with commas and no intervening spaces. This parameter is optional.

Admin

Specifies the name of the administrator whose password expiration period you would like to reset. To specify a list of administrators, separate the names with commas and no intervening spaces. This parameter is optional.

Example: Reset the password expiration for specific client nodes

Reset the password expiration period for client nodes bj and katie.

```
reset passexp node=bj,katie
```

Example: Reset the password expiration for all users

Reset the password expiration period for all users to the common expiration period.

```
reset passexp
```

Related commands

Table 1. Commands related to RESET PASSEXP

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes that are associated with a client node.

RESTART EXPORT (Restart a suspended export operation)

Use this command to restart a suspended export operation.

An export operation is suspended when any of the following conditions is detected:

- A SUSPEND EXPORT command is issued for the running export operation
- Segment preemption - the file being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered

Important: Nodes or file spaces (on the exporting server) in the original export operation that are subsequently renamed are not included in the resumed operation. Any remaining data for nodes or file spaces on the target server that are deleted prior to resumption are discarded.

Privilege class

You must have system privilege to issue this command.

Syntax

```
>>-RESTART EXPORT .-*-----
                    +-----+-----><
                    '---export_identifier---
```

Parameters

export_identifier

This optional parameter is the unique identifier for the suspended server-to-server export operation. You can use the wildcard character to specify this name. The export identifier name can be found by issuing the QUERY EXPORT command to list all the currently suspended server-to-server export operations.

Example: Restart a suspended export

Restart the suspended export operation identified by the export identifier EXPORTALLACCTNODES.

```
restart export exportallacctnodes
```

Related commands

Table 1. Commands related to RESTART EXPORT

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.

Command	Description
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
SUSPEND EXPORT	Suspends a running export operation.

RESTORE commands

Use the RESTORE commands to restore IBM Spectrum Protect™ storage pools or volumes.

- RESTORE NODE (Restore a NAS node)
- RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)
- RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)

RESTORE NODE (Restore a NAS node)

Use this command to initiate a restore operation for a network-attached storage (NAS) node.

You can use the RESTORE NODE command to restore backups that were created by using either the client's BACKUP NAS command or the server's BACKUP NODE command. NAS data may be restored from primary or copy native IBM Spectrum Protect™ pools; primary or copy NAS pools; or any combination needed to achieve the restore.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax

```
>>-RESTORE Node--node_name--source_file_system----->
    .-source_file_system-----
>--+-----+----->
    '-destination_file_system-'

>--+-----+----->
    |                                     |
    |         .-,----- .             |
    |         v                       |
    '-FILELIST--==+---file_name+---+-'
        '-FILE:--file_list-'

    .-NAMEType--==--SERVER----- .
>--+-----+----->
    '-NAMEType--==--+SERVER-----+-'
        +-HEXadecimal-+
        '-UNICODE-----'

    .-PITDate--==--TODAY----- .
>--+-----+----->
    '-PITDate--==--+mm/dd/yyyy+---+-'
        +-TODAY-----+
        +-TODAY-numdays+
        '- -numdays-----'

    .-PITTime--==--NOW----- .    .-Wait--==--No----- .
>--+-----+-----+----->
    '-PITTime--==--+hh:mm:ss--+-'    '-Wait--==--+No--+-'
        +-NOW-----+                '-Yes-'
        +-NOW-hh:mm-+

```


SERVER

The server uses the server's code page to interpret the names.

HEXadecimal

The server interprets the names that you enter as the hexadecimal representation of a name in Unicode. To view the hexadecimal representation of a file or directory name, you can use the QUERY TOC command with FORMAT=DETAILED.

UNIcode

The server interprets the names as being UTF-8 encoded. This option only applies when you have specified a list with FILELIST=FILE:file_list.

Restriction: Network Data Management Protocol (NDMP) has limitations that prevent IBM Spectrum Protect from reporting whether or not individual files and directories are successfully restored.

PITDate

Specifies the point-in-time date. When used with the PITTIME parameter, PITDATE establishes the point in time from which you want to select the data to restore. The latest data that was backed up on or before the date and time that you specify will be restored. This parameter is optional. The default is TODAY.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	06/25/2001
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified	TODAY-7 or -7. To restore data that was backed up a week ago, specify PITDATE=TODAY-7 or PITDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

PITTime

Specifies the point-in-time time. When used with the PITDATE parameter, PITTIME establishes the point in time from which you want to select the data to restore. The latest data that was backed up on or before the date and time that you specify will be restored. This parameter is optional. The default is the current time.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified date	12:33:28
NOW	The current time on the specified date	NOW
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue this command at 9:00 with PITTIME=NOW-03:30 or PITTIME=-03:30, the server restores backup records with a time of 5:30 or later on the point-in-time date.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is NO.

Possible values are:

No

Specifies that the server processes this command in the background. Use the QUERY PROCESS command to monitor the background processing of this command.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

TYPE

Specifies the type of image to restore. The default value for this parameter is BACKUPIIMAGE and it is used to restore data from standard NDMP base or differential backups. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

BACKUPIImage

Specifies that the file system should be restored from the appropriate standard NDMP backup images. This is the default method for performing an NDMP restore operation. Using the BACKUPIIMAGE type, you can restore data from base and differential backups, and data at the file level.

SNAPMirror

Specifies that the file system should be retrieved from a NetApp SnapMirror image. SnapMirror images are block-level full-backup images of a NetApp file system. A SnapMirror image can only be restored to a file system that has been prepared as a SnapMirror target volume. Refer to the documentation that came with your NetApp file server for details.

After a SnapMirror image is retrieved and copied to a target file system, IBM Spectrum Protect breaks the SnapMirror relationship that was created by the file server during the operation. After the restore is complete, the target file system returns to the same state as that of the original file system at the point-in-time of the backup.

When setting the TYPE parameter to SNAPMIRROR, note the following restrictions:

Restrictions:

- You cannot specify the FILELIST parameter.
- Neither the *source_file_system_name* nor the *destination_file_system_name* can be a virtual filespace name.
- This parameter is valid for NetApp and IBM® N-Series file servers only.

Example: Restore a complete directory

Restore all of the files and subdirectories in the directory `/mydir`.

```
restore node nasnode /myfs /dest filelist=/path/to/mydir
```

Example: Restore data from a file system

Restore the data from the `/vol/vol10` file system on NAS node `NAS1`.

```
restore node nas1 /vol/vol10
```

Example: Restore a directory-level backup to the same location

Restore the directory-level backup to the original location. The source is the virtual file space name `/MIKESDIR` and no destination is specified.

```
restore node nas1 /mikesdir
```

For this example and the next example, assume the following virtual file space definitions exist on the server for the node `NAS1`.

VFS Name	Filesystem	Path
<code>/mikesdir</code>	<code>/vol/vol2</code>	<code>/mikes</code>
<code>/TargetDirVol2</code>	<code>/vol/vol2</code>	<code>/tmp</code>
<code>/TargetDirVol1</code>	<code>/vol/vol1</code>	<code>/tmp</code>

Example: Restore a directory-level backup to a different file system

Restore the directory-level backup to a different file system but preserve the path.

```
restore node nas1 /mikesdir /vol/vol0
```

Related commands

Table 1. Commands related to RESTORE NODE

Command	Description
BACKUP NODE	Backs up a network-attached storage (NAS) node.
CANCEL PROCESS	Cancels a background server process.
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY NASBACKUP	Displays information about NAS backup images.
QUERY TOC	Displays details about the table of contents for a specified backup image.

RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)

Use this command to restore files from one or more copy storage pools or active-data pools to a primary storage pool.

IBM Spectrum Protect™ restores all the primary storage pool files that:

- Have been identified as having errors
- Reside on a volume with an access mode of DESTROYED

Restriction: You cannot use this command for container storage pools. Use the REPLICATE STGPOOL command to protect data for container storage pools.

You can also use this command to identify volumes that contain damaged, primary files. During restore processing, a message is issued for every volume in the restored storage pool that contains damaged, non-cached files. Use the QUERY CONTENT command to identify damaged, primary files on a specific volume.

You cannot restore a storage pool defined with a CENTERA device class.

In addition to restoring data to primary storage pools that have NATIVE or NONBLOCK data formats, this command also lets you restore data to primary storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The primary storage pool must have the same data format as the copy storage pool from which data is to be restored. IBM Spectrum Protect supports backend data movement for NDMP images.

Tip: To restore NAS client-node data to NAS storage pools, you must manually change the access mode of the volumes to DESTROYED using the UPDATE VOLUME command. However, if you are using disaster recovery manager, the plan file will contain the information the server needs to automatically mark the volumes as DESTROYED.

Restoration of files might be incomplete if backup file copies in copy storage pools or active-data pools were moved or deleted by other IBM Spectrum Protect processes during restore processing. To prevent this problem, do not issue the following commands for copy storage pool or active-data pool volumes while restore processing is in progress:

- MOVE DATA
- DELETE VOLUME (DISCARDATA=YES)
- AUDIT VOLUME (FIX=YES)

Also, you can prevent reclamation processing for your copy storage pools by setting the RECLAIM percentage to 100 with the UPDATE STGPOOL command.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool for which files are to be restored. If you are a restricted storage administrator and you want to restore files to a new primary storage pool, you must also have authority for the new storage pool.

Syntax

```
>>-RESTORE STGpool--primary_pool_name----->
>+-----+
  '-COPYstgpool----copy_pool_name-'
```

```

.-ACTIVEDATAOnly---No-----
>--+-----+----->
'-ACTIVEDATAOnly---+No-----+'
          '-Yes--| A |- '

>--+-----+----->
'-NEWstgpool-----new_primary_pool_name-'

.-MAXPRocess-----1----- .-Preview-----No-----
>--+-----+-----+----->
'-MAXPRocess-----number-' '-Preview-----+No--+-'
                                '-Yes-'

.-Wait-----No-----
>--+-----+-----><
'-Wait-----+No--+-'
          '-Yes-'

A (Yes)

|--ACTIVEDATAPool-----active-data_pool_name-----|

```

Parameters

primary_pool_name (Required)

Specifies the name of the primary storage pool that is being restored.

COPYstgpool

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any copy pool in which copies can be located. Do not use this parameter with the `ACTIVEDATAONLY` or `ACTIVEDATAPool` parameters.

ACTIVEDATAOnly

Specifies that active versions of backup files are to be restored from active-data pools only. This parameter is optional. The default is `NO`. If this parameter is not specified, files are restored from copy-storage pools. Do not use this parameter with the `COPYSTGPPOOL` parameter. Possible values are:

No

Specifies that the storage pool will not be restored from active-data pools.

Yes

Specifies that the storage pool will be restored from active-pool or pools that you specify using the `ACTIVEDATAPool` parameter. If you specify `YES` as a value for `ACTIVEDATAONLY`, but do not specify a value for `ACTIVEDATAPool`, files are restored from any active-data pool in which active versions of backup files can be located.

Attention: Restoring a primary storage pool from an active-data pool might cause some or all inactive files to be deleted from the database if the server determines that an inactive file needs to be replaced but cannot find it in the active-data pool.

ACTIVEDATAPool

Specifies the name of the active-data pool from which the active versions of backup files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any active-data pool in which active versions of backup files can be located.

NEWstgpool

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If this parameter is not specified, files are restored to the original primary storage pool (the pool being restored).

MAXPRocess

Specifies the maximum number of parallel processes that are used for restoring files. Using multiple, parallel processes may improve throughput for the restore. This parameter is optional. You can specify a value from 1 to 999. The default is 1.

When determining this value, consider the number of mount points (logical drives) and physical drives that can be dedicated to this operation. To access a sequential access volume, IBM Spectrum Protect uses a mount point, and, if the device type is not `FILE`, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the restore.

Each process needs a mount point for copy storage pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are restoring files in a sequential storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device class is not FILE, an additional drive. For example, suppose you specify a maximum of 3 processes to restore a primary sequential storage pool from a copy storage pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least 6, and at least 6 mount points and 6 drives must be available.

To preview a restore, only one process is used and no mount points or drives are needed.

Preview

Specifies if you want to preview but not perform the restore. The preview lets you identify volumes required to restore the storage pool. The preview displays:

- A list of primary storage pool volumes that contain damaged files.
- The number of files and the number of bytes to be restored, assuming that the access mode of the required copy storage pool volumes is READWRITE or READONLY when the restore operation is performed.
- A list of copy storage pool volumes containing files to be restored. These volumes must be mounted if you perform the restore.
- A list of any volumes containing files that cannot be restored.

Note: For only a list of offsite copy storage pool volumes to be mounted during a restore, change the access mode of the copy pool volumes to UNAVAILABLE. This prevents reclamation and move data processing of the volumes until they are moved onsite for the restore.

This parameter is optional. The default is NO. Possible values are:

No

Specifies that the restore is done.

Yes

Specifies that you want to preview the restore but not do the restore.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed.

Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged. To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files may have already been restored prior to the cancellation.

Yes

Specifies that the server performs this operation in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the operation completes.

Note: You cannot specify WAIT=YES from the server console.

Example: Restore files from a copy storage pool to the primary storage pool

Restore files from any copy storage pool to the primary storage pool, PRIMARY_POOL.

```
restore stgpool primary_pool
```

Example: Restore files from a specific active-data pool to the primary storage pool

Restore files from active-data pool ADP1 to the primary storage pool PRIMARY_POOL.

```
restore stgpool primary_pool activedataonly=yes activedatapool=adp1
```

Related commands

Table 1. Commands related to RESTORE STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)

Use this command to restore all files on damaged volumes in a primary storage pool that was backed up to a copy storage pool or copied to an active-data pool. IBM Spectrum Protect™ does not restore cached copies of files and removes those cached files from the database during restore processing.

In addition to restoring data to volumes in storage pools that have NATIVE or NONBLOCK data formats, this command also lets you restore data to volumes in storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The volumes to be restored must have the same data format as the volumes in the copy storage pool. IBM Spectrum Protect supports backend data movement for NDMP images.

This command changes the access mode of the specified volumes to DESTROYED. When all files on a volume are restored to other locations, the destroyed volume is empty and is deleted from the database.

The restoration may be incomplete for one or more of the following reasons:

- Files were either never backed up or the backup copies are marked as damaged. Use the QUERY CONTENT command to get more information on the remaining files on the volume.
- A copy storage pool was specified on the RESTORE command, but files were backed up to a different copy storage pool. Use the PREVIEW parameter when you issue the RESTORE command again to determine if this is the problem.
- Volumes in the copy storage pool needed to perform the restore operation are offsite or unavailable. Check the activity log for messages that occurred during restore processing.
- Backup file copies in copy storage pools were moved or deleted by other processes during a restore. See note 3.
- An active-data pool was specified for the restore, and inactive files were not available to be copied.

Important:

1. You cannot restore volumes in storage pools defined with a CENTERA device class.
2. Before you restore a random-access volume, issue the VARY command to vary the volume offline.
3. To prevent copy storage pools files from being moved or deleted by other processes, do not issue the following commands for copy storage pool volumes during a restore:
 - MOVE DATA
 - DELETE VOLUME (DISCARDATA=YES)
 - AUDIT VOLUME (FIX=YES)

To prevent reclamation processing of copy storage pools, issue the UPDATE STGPOOL command with the RECLAIM parameter set to 100.

Privilege class

To issue this command you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool. If you have restricted privilege and want to restore files to a new primary storage pool, you must also have authority for the new storage pool.

Syntax

```

      .- ,----- .
      V           |
>>-RESTORE Volume---volume_name+----->
>--+-----+----->
' -COPYstgpool----copy_pool_name-'
      .-ACTIVEDATAOnly---No----- .
>--+-----+----->
' -ACTIVEDATAOnly---+No-----+'
      '-Yes--| A |-'
>--+-----+----->
' -NEWstgpool----new_primary_pool_name-'
      .-MAXPRocess----1----- .   .-Preview----No----- .
>--+-----+----->
' -MAXPRocess----number-'   '-Preview----+No---+'
      '-Yes-'
      .-Wait----No----- .
>--+-----+-----><
' -Wait----+No---+'
      '-Yes-'

A (Yes)

|--ACTIVEDATAPool----active-data_pool_name-----|

```

Parameters

volume_name (Required)

Specifies the name of the primary storage pool volume to be restored. To specify a list of volumes that belong to the same primary storage pool, separate the names with commas and no intervening spaces.

COPYstgpool

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If you do not specify this parameter, files are restored from any copy pool in which copies can be located. Do not use this parameter with the `ACTIVEDATAONLY` or `ACTIVEDATAPOOL` parameters.

ACTIVEDATAOnly

Specifies that active versions of backup files are to be restored from active-data pools only. This parameter is optional. The default is `NO`. If this parameter is not specified, files are restored from copy-storage pools. Do not use this parameter with the `COPYSTGPPOOL` parameter. Possible values are:

No

Specifies that the storage pool will not be restored from active-data pools.

Yes

Specifies that the storage pool will be restored from active-pool or pools that you specify using the `ACTIVEDATAPOOL` parameter. If you specify `YES` as a value for `ACTIVEDATAONLY`, but do not specify a value for `ACTIVEDATAPOOL`, files are restored from any active-data pool in which active versions of backup files can be located.

Attention: Restoring a volume from an active-data pool might cause some or all inactive files to be deleted from the database if the server determines that an inactive file needs to be replaced but cannot find it in the active-data pool.

ACTIVEDATAPool

Specifies the name of the active-data pool from which the active versions of backup files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any active-data pool in which active versions of backup files can be located.

NEWstgpool

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If you do not specify this parameter, files are restored to the original primary storage pool.

MAXPRocess

Specifies the maximum number of parallel processes to use for restoring files. Using parallel processes may improve throughput. This parameter is optional. You can specify a value from 1 to 999. The default is 1.

When determining this value, consider the number of mount points (logical drives) and physical drives that can be dedicated to this operation. To access a sequential access volume, IBM Spectrum Protect uses a mount point, and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the restore.

Each process needs a mount point for copy storage pool volumes. If the device type is not FILE, each process also needs a drive. If you are restoring a sequential storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device type is not FILE, an additional drive. For example, suppose you specify a maximum of three processes to back up a primary sequential storage pool to a copy storage pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least 6, and at least 6 mount points and 6 drives must be available.

To preview a backup, only one process is used and no mount points or drives are needed.

Preview

Specifies if you want to preview but not perform the restore. You can use this option to identify the offsite volumes required to restore a storage pool. This parameter is optional. The default is NO. Possible values are:

No

Specifies that you want to perform the restore operation.

Yes

Specifies that you want to preview the restore operation but restore the data.

Tip: If you preview a restore to see a list of offsite copy pool volumes to be mounted, you should change the access mode of the identified volumes to UNAVAILABLE. This prevents reclamation and MOVE DATA processing for these volumes until they are transported to the onsite location for use in restore processing.

The preview displays the following:

- The number of files and bytes to be restored, if the access mode of the copy storage pool volumes is READWRITE or READONLY when the restoration is performed.
- A list of copy storage pool volumes containing files to be restored. These volumes must be mounted if you perform the restore.
- A list of volumes containing files that cannot be restored.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. This default is NO. Possible values are:

No

Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If you cancel this process, some files may have already been backed up prior to the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the command completes.

Remember: You cannot specify WAIT=YES from the server console.

Example: Restore primary volume data files

Restore files stored on volume PVOL2 in primary storage pool PRIMARY_POOL.

```
restore volume pvol2
```

Example: Restore primary volume data files from an active-data pool

Restore files stored on volume VOL001 in primary pool PRIMARY_POOL from active-data pool ADP1.

```
restore volume vol001 activedataonly=yes activedatapool=adp1
```


Authority for all matching storage pools will be revoked. If STGPOOLS is specified then the parameter CLASSES=STORAGE is optional.

Usage notes

1. To change an unrestricted storage administrator to a restricted storage administrator, you must first use this command to revoke the unrestricted privilege. Then, use the GRANT AUTHORITY command to grant the administrator restricted storage privilege and to identify the storage pools to which the administrator has authority.

To revoke unrestricted storage privilege from an administrator, specify the CLASSES=STORAGE parameter. You cannot use the STGPOOLS parameter to revoke authority for selected storage pools from an unrestricted storage administrator.

2. To change an unrestricted policy administrator to a restricted policy administrator, you must first use this command to revoke the unrestricted privilege. Then, use the GRANT AUTHORITY command to grant the administrator restricted policy privilege and to identify the policy domains to which the administrator has authority.

To revoke unrestricted policy privilege from an administrator, specify the CLASSES=POLICY parameter. You cannot use the DOMAINS parameter to revoke authority for selected domains from an unrestricted administrator.

Example: Revoke certain administrative privileges

Revoke part of administrator CLAUDIA's privileges. CLAUDIA has restricted policy privilege for the policy domains EMPLOYEE_RECORDS and PROG1. Restrict CLAUDIA's policy privilege to the EMPLOYEE_RECORDS policy domain.

```
revoke authority claudia classes=policy
domains=employee_records
```

Example: Revoke all administrative privileges

Administrator LARRY currently has operator and restricted policy privilege. Revoke all administrative privileges for administrator LARRY. To revoke all administrative privileges for an administrator, identify the administrator, but do not specify CLASSES, DOMAINS, or STGPOOLS. LARRY remains an administrator but he can only use those commands that can be issued by any administrator.

```
revoke authority larry
```

Example: Revoke node privilege

Help desk personnel user CONNIE currently has node privilege with client owner authority for client node WARD3. Revoke her node privilege with client owner authority.

```
revoke authority connie classes=node
authority=owner node=ward3
```

Related commands

Table 1. Commands related to REVOKE AUTHORITY

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect/IBM Spectrum Protect™ administrators.

REVOKE PROXYNODE (Revoke proxy authority for a client node)

Use this command to revoke authority for an agent client node to perform backup and restore operations for a target node on the IBM Spectrum Protect™ server.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege

Syntax

```
>>-REvOke PROXynode TArget-----target_node_name----->
>--AGent-----agent_node_name-----<<
```

Parameters

TArget (Required)

Specifies the target node to which an agent node has been granted proxy authority. Wildcard characters and comma-separated lists of node names are allowed.

AGent (Required)

Specifies which node has authority to act as proxy to the target node. Wildcard characters and comma-separated lists of node names are allowed.

Example: Revoke a node's proxy authority

To revoke authority from target node NASCLUSTER to act as proxy for all agent nodes which start with the letter M, issue the following command.

```
revoke proxynode target=nascluster agent=m*
```

Related commands

Table 1. Commands related to REVOKE PROXYNODE

Command	Description
GRANT PROXYNODE	Grant proxy authority to an agent node.
QUERY PROXYNODE	Display nodes with authority to act as proxy nodes.

ROLLBACK (Rollback uncommitted changes in a macro)

Use this command within a macro to undo any processing changes made by commands run by the server but not yet committed to the database. A committed change is permanent and cannot be rolled back. The ROLLBACK command is useful for testing macros.

Ensure that your administrative client session is not running with the ITEMCOMMIT option when using this command.

Important: SETOPT commands inside a macro cannot be rolled back.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-ROLLBACK-----<<
```

Parameters

None

Example: Rollback changes in a macro

Run the REGN macro with the ROLLBACK command to verify that the macro works without committing any changes. The macro contents are:


```

/* Macro to register policy
administrators and grant authority */
REGister Admin sara hobby
GRant AUTHority sara CClasses=Policy
REGister Admin ken plane
GRant AUTHority ken CClasses=Policy
ROLLBACK /* prevents any changes from being committed */

```

Related commands

Table 1. Commands related to ROLLBACK

Command	Description
COMMIT	Makes changes to the database permanent.
MACRO	Runs a specified macro file.

Related concepts:

Administrative client macros

RUN (Run an IBM Spectrum Protect script)

Use this command to run an IBM Spectrum Protect™ script. To issue this command on another server, the script being run must be defined on that server.

You can include RUN commands in scripts as long as they do not create loops. For example, you should avoid including RUN commands where SCRIPT_A runs SCRIPT_B and SCRIPT_B runs SCRIPT_A.

Important: IBM Spectrum Protect does not have a command that can cancel a script after it starts. To stop a script, you must halt the server.

Privilege class

To issue this command, you must have operator, policy, system, storage, or system privilege.

Syntax

```

>>-RUN--script_name--+-----+----->
      | .-,------. |
      | v             | |
      +---substitution_value+--'

.-Preview---No-----.-Verbose---No-----
>--+-----+----->>
'-Preview---+No--+-' '-Verbose---+No--+-'
      '-Yes-'           '-Yes-'

```

Parameters

script_name (Required)

Specifies the name of the script you want processed. The name you specify cannot be a substitution variable, such as \$1.

substitution_value

Specifies one or more values to substitute for variables when the script is run. In a script, a substitution variable consists of a '\$' character, followed by a number. When you run the script, IBM Spectrum Protect replaces the substitution variables defined in a script with the values you supply with this command. You must specify values for each substitution variable defined in the script or the script will fail. This parameter is optional.

Preview

Specifies whether to preview the command lines of a script without actually processing the script. The default is NO. Possible values are:

Yes

Specifies that the command lines included in a script are displayed, but the script is not processed.

No

Specifies that the command lines included in a script are displayed and the script is processed.

Verbose

Specifies whether command lines, variable substitution, and conditional logic testing used in a script are displayed as the script is being processed. This parameter is ignored if PREVIEW=YES is specified. The default is NO.

Possible values are:

Yes

Specifies that the command lines, variable substitution, and conditional logic testing are displayed as the script is being processed.

No

Specifies that the command lines, variable substitution, and conditional logic testing do not display as the script is being processed.

Example: View the commands generated by a script with a table name substitution variable

To run the following example script, called QSAMPLE, you issue a RUN command that specifies the table name ACTLOG as the value for the substitution variable, \$1. Use the output to preview the commands generated by the script before running the commands.

```
001 /* This is a sample SQL Query in wide format */
005 SET SQLDISPLAYMODE WIDE
010 SELECT colname FROM -
015 COLUMNS WHERE TABNAME='$1'

run qsample actlog preview=yes

ANR1461I RUN: Executing command script QSAMPLE.
ANR1466I RUN: Command script QSAMPLE, Line 5 :
           set sqldisplaymode wide.
ANR1466I RUN: Command script QSAMPLE, Line 15 :
           select colname from columns where tabname='ACTLOG'.
ANR1470I RUN: Command script QSAMPLE completed successfully
           (PREVIEW mode)
```

Example: Run a script to display and run the commands generated by the script

Run the same script as show in the prior example to display both the generated commands and the results of the commands.

```
run qsample actlog verbose=yes

ANR1461I RUN: Executing command script QSAMPLE.
ANR1466I RUN: Command script QSAMPLE, Line 5 :
           set sqldisplaymode wide.
ANR1466I RUN: Command script QSAMPLE, Line 5 : RC=RC_OK
ANR1466I RUN: Command script QSAMPLE, Line 15 :
           select colname from columns where tabname='ACTLOG'.

COLNAME
-----
DATE_TIME
MSGNO
SEVERITY
MESSAGE
ORIGINATOR
NODENAME
OWNERNAME
SCHEDNAME
DOMAINNAME
SESSID

ANR1462I RUN: Command script QSAMPLE, Line 15 : RC=RC_OK
ANR1462I RUN: Command script QSAMPLE completed successfully.
```

Example: Run a script to display just the results of the commands in the script

Run the previous example script, without displaying just the results of the generated commands in the script.

```
run qsample actlog verbose=no

COLNAME
-----
```

DATE_TIME
MSGNO
SEVERITY
MESSAGE
ORIGINATOR
NODENAME
OWNERNAME
SCHEDNAME
DOMAINNAME
SESSID

ANR1462I RUN: Command script QSAMPLE completed successfully.

Related commands

Table 1. Commands related to RUN

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
UPDATE SCRIPT	Changes or adds lines to a script.

Related tasks:

Running a server script

SELECT (Perform an SQL query of the IBM Spectrum Protect database)

Use the SELECT command to create and format a customized query of the IBM Spectrum Protect™ database.

IBM Spectrum Protect provides an SQL interface to a DB2® program. Restrictions and guidelines for handling SQL queries are handled directly by DB2.

To help you find what information is available, IBM Spectrum Protect provides three system catalog tables:

SYSCAT.TABLES

Contains information about all tables that can be queried with the SELECT command.

SYSCAT.COLUMNS

Describes the columns in each table.

You can issue the SELECT command to query these tables to determine the location of the information that you want.

Usage notes

You cannot issue the SELECT command from a server console.

Because the select command does not lock and unlock records, contention for a record can cause the server to erroneously issue message ANR2034E: *SELECT: No match found using this criteria*. Check your selection criteria, and if you believe that it is correct, try the command again.

To stop the processing of a SELECT command after it starts, cancel the administrative session from which the command was issued. Cancel the session from either the server console or another administrative session.

Temporary table spaces are used to process SQL queries within DB2. Inadequate temporary space can cause SQL queries to fail.

To export output to a comma-separated file for import into a spreadsheet, use -comma and > command-line options on the dsmdmc command.

Privilege class

Any administrator can issue this command.

Syntax

For SELECT statement syntax and guidelines, search the DB2 product information.

Important: The appropriate syntax for the timestamp Select statement is:

```
SELECT * FROM SUMMARY WHERE ACTIVITY='EXPIRATION' AND START_TIME >'2009-05-10 00:00:00' AND
START_TIME <'2009-05-11 23:23:23'
```

List of examples

The SELECT command is used to customize a wide variety of queries. To give you an idea of what you can do with the command, this section includes many examples. There are, however, many more possibilities. Query output is shown only for the more complex commands to illustrate formatting.

The following list summarizes the example SELECT commands:

- List administrator user ID passwords that are authenticated with an external LDAP directory server
- List available tables
- List client nodes and administrative clients that are currently locked from server access
- List client nodes and administrative clients that have not specified the correct password lately
- List nodes in the standard policy domain that are not associated with the daily backup schedule DAILYBACKUP
- List the administrators that have policy authority
- List type E (ERROR) or W (WARNING) messages that have been issued in the time period for which activity log records have been maintained
- List the administrative schedules that have been defined or altered by administrator JAKE
- List the relative administrative schedule priorities
- List the management classes that have an archive copy group with a retention period greater than 365 days
- List the client nodes that are in each policy domain
- Count how many files have been archived from each node
- List the clients that are using space management
- Determine how many volumes would be reclaimed if the reclamation threshold is changed to 50 percent for storage pool TAPE
- Determine how many backup files would be affected for each node if the DAILY management class in the STANDARD policy domain is changed or deleted
- For all active client sessions, determine how long have they been connected and their effective throughput in bytes per second
- Determine how long the current background processes have been running and determine their effective throughput in time and files per second
- Count the number of client nodes are there for each platform type
- Count the number of file spaces each client node has and list the client nodes ascending order
- Obtain statistical information for calculating the number of off-site volumes that have their space reclaimed during reclamation of a storage pool
- Obtain PVU estimate detail records
- Obtain information about the node roles
- Obtain information about status

Example: List administrator user IDs that authenticate to the IBM Spectrum Protect server

List all the administrator user IDs whose passwords authenticate with the IBM Spectrum Protect server:

```
select admin_name from admins where
authentication=local
```

Example: List available tables

List all the tables available for querying the IBM Spectrum Protect database.

```
select * from syscat.tables

      ABSHEMA: SERVER1
      TABNAME: ACTLOG
CREATE_TIME: 1999-05-01 07:39:06
  COLCOUNT: 10
INDEX_COLCOUNT: 1
```

```

UNIQUE_INDEX: FALSE
REMARKS: Server activity log

TABSHEMA: SERVER1
TABNAME: ADMIN_SCHEDULES
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 14
INDEX_COLCOUNT: 1
UNIQUE_INDEX: TRUE
REMARKS: Administrative command schedules

TABSHEMA: SERVER1
TABNAME: ADMINS
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 15
INDEX_COLCOUNT: 1
UNIQUE_INDEX: TRUE
REMARKS: Server administrators

TABSHEMA: SERVER1
TABNAME: ARCHIVES
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 10
INDEX_COLCOUNT: 5
UNIQUE_INDEX: FALSE
REMARKS: Client archive files

```

Example: List client nodes and administrative clients that are currently locked from server access

```

select node_name from nodes where locked='YES'

select admin_name from admins where locked='YES'

```

Example: List client nodes, administrative clients, and servers that are using transitional session security

```

select node_name from nodes where session_security='Transitional'

select admin_name from admins where session_security='Transitional'

select server_name from servers where session_security='Transitional'

```

Example: List client nodes and administrative clients that have not specified the correct password lately

```

select node_name from nodes where invalid_pw_count <>0

select admin_name from admins where invalid_pw_count <>0

```

Example: List nodes in the standard policy domain that are not associated with the daily backup schedule DAILYBACKUP

```

select node_name from nodes where domain_name='STANDARD' and
node_name not in (select node_name from associations
where domain_name='STANDARD' and
schedule_name='DAILYBACKUP')

```

Example: List the administrators who have policy authority

```

select admin_name from admins where
upper(system_priv) <>'NO'
or upper(policy_priv) <>'NO'

```

Example: List type E (ERROR) or W (WARNING) messages that have been issued in the time period for which activity log records have been maintained

```
select date_time,msgno,message from actlog
where severity='E' or severity='W'
```

Example: List the administrative schedules that have been defined or altered by administrator JAKE

```
select schedule_name from admin_schedules
where chg_admin='JAKE'
```

Example: List the relative administrative schedule priorities

```
select schedule_name,priority from admin_schedules order
by priority
```

Example: List the management classes that have an archive copy group with a retention period greater than 365 days

```
select domain_name,set_name,class_name from ar_copygroups
where retver='NOLIMIT' or cast(retver as integer) >365
```

Example: List the management classes that specify more than five backup versions

```
select domain_name,set_name,class_name from bu_copygroups
where verexists ='NOLIMIT' or
cast(verexists as integer)>5
```

Example: List the client nodes that are using the client option set named SECURE

```
select node_name from nodes where option_set='SECURE'
```

Example: List the client nodes that are in each policy domain

```
select domain_name,num_nodes from domains
```

Example: Count how many files have been archived from each node

Attention: This command might take a long time to complete.

```
select node_name,count(*) from archives
group by node_name
```

Example: List the clients that are using space management

```
select node_name from auditocc where spacemg_mb <>0
```

Example: Determine how many volumes would be reclaimed if the reclamation threshold is changed to 50 percent for storage pool TAPE

```
select count(*) from volumes where stgpool_name='TAPE'
and upper(status)='FULL' and pct_utilized < 50
```

Example: Determine how many backup files would be affected for each node if the DAILY management class in the STANDARD policy domain is changed or deleted

Note: This command takes significant time and resources to complete.

```
select node_name, count(*) as "Files" from backups
where class_name='DAILY' and node_name in
(select node_name from nodes where domain_name='STANDARD')
group by node_name
```

Example: For all active client sessions, determine how long have they been connected and their effective throughput in bytes per second

```

select session_id as "Session",
client_name as "Client",
state as "State",
current_timestamp-start_time as "Elapsed Time",
(cast(bytes_sent as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes sent/second",
(cast(bytes_received as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes received/second"
from sessions

```

```

      Session: 24
      Client: ALBERT
      State: Run
      Elapsed Time: 0 01:14:05.000000
      Bytes sent/second: 564321.9302768451
      Bytes received/second: 0.0026748857944

```

```

      Session: 26
      Client: MILTON
      State: Run
      Elapsed Time: 0 00:06:13.000000
      Bytes sent/second: 1638.5284210992221
      Bytes received/second: 675821.6888561849

```

Example: Determine how long the current background processes have been running and determine their effective throughput in time and files per second

Note: Expiration does not report the number of bytes processed.

```

select process_num as "Number",
process,
current_timestamp-start_time as "Elapsed Time",
(cast(files_processed as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Files/second",
(cast(bytes_processed as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes/second"
from processes

```

```

      Number: 1
      PROCESS: Expiration
      Elapsed Time: 0 00:24:36.000000
      Files/second: 6.3216755870092
      Bytes/second: 0.000000000000000

```

Example: Count the number of client nodes for each platform type

```

select platform_name,count(*) as "Number of Nodes"
from nodes group by platform_name

```

PLATFORM_NAME	Number of Nodes
AIX	6
SunOS	27
Win32	14
Linux	20

Example: Count the number of file spaces each client node has and list the client nodes ascending order

```

select node_name, count(*) as "number of filespaces"
from filespaces group by node_name order by 2

```

NODE_NAME	number of filespaces
-----	-----

ALBERT	2
MILTON	2
BARNEY	3
SEBASTIAN	3
MAILHOST	4
FALCON	4
WILBER	4
NEWTON	4
JEREMY	4
WATSON	5
RUSSELL	5

Example: Obtain statistical information for calculating the number of off-site volumes that have their space reclaimed during reclamation of a storage pool.

```
select * from summary where activity='OFFSITE RECLAMATION'

START_TIME: 2004-06-16 13:47:31.000000
END_TIME: 2004-06-16 13:47:34.000000
ACTIVITY: OFFSITE RECLAMATION
NUMBER: 4
ENTITY: COPYPOOL
COMMMETH:
ADDRESS:
SCHEDULE_NAME:
EXAMINED: 170
AFFECTED: 170
FAILED: 0
BYTES: 17821251
IDLE: 0
MEDIAP: 0
PROCESSES: 2
SUCCESSFUL: YES
VOLUME_NAME:
DRIVE_NAME:
LIBRARY_NAME:
LAST_USE:
COMM_WAIT:
NUM_OFFSITE_VOLS: 2
```

Example: Identify which storage pools contain data that was deduplicated by clients

```
select stgpool_name,has_client_dedup_data from stgpools

STGPOOL_NAME          HAS_CLIENT_DEDUP_DATA
-----
ADPOOL                NO
ARCHIVEPOOL           NO
BACKUPPOOL            NO
COPYDEDUP             NO
COPYNODEDUP          NO
FILEPOOL              YES
FILEPOOL2             NO
LANFREEFILEPOOL      YES
SPACEMGPOOL          NO
```

Example: Obtain information about the database

```
select * from db

DATABASE_NAME: TSMDB1
TOT FILE SYSTEM MB: 2048000
USED_DB_SPACE_MB: 12576
FREE_SPACE_MB: 1576871
TOTAL_PAGES: 983044
USABLE_PAGES: 982908
USED_PAGES: 977736
FREE_PAGES: 5172
BUFF_HIT_RATIO: 96.2
TOTAL_BUFF_REQ: 53967
SORT_OVERFLOW: 0
```



```

LOCK_ESCALATION: 0
PKG_HIT_RATIO: 70.0
LAST_REORG: 2010-07-15 17:32:55.000000
FULL_DEV_CLASS: OUTFILE
NUM_BACKUP_INCR: 0
LAST_BACKUP_DATE: 2010-01-21 10:37:59.000000
PHYSICAL_VOLUMES: 0
PAGE_SIZE:
NUM_BACKUP_STREAMS: 4

```

Example: Obtain PVU estimate detail records

Generate the PVU estimate for a node named ACCTSRECSRV, which is used by the IBM Spectrum Protect Extended Edition product.

```
select * from pvuestimate_details where node_name='ACCTSRECSRV'
```

```

PRODUCT: PRODEE
LICENSE_NAME: MGSYSLAN
NODE_NAME: ACCTSRECSRV
LAST_USED: 2008-01-20 16:12:24.000000
TRYBUY: FALSE
PROC_VENDOR: IBM
PROC_BRAND: POWER5+ QCM
PROC_TYPE: 4
PROC_MODEL:
PROC_COUNT: 2
ROLE: SERVER
ROLE_OVERRIDE: USEREPORTED
ROLE_EFFECTIVE: SERVER
VALUE_UNITS: 50
VALUE_FROM_TABLE: YES
PVU: 100
SCAN_ERROR : NO
API_CLIENT: NO
PVU_AGNOSTIC: NO
HYPERVISOR: VMWARE
GUID: 01.2e.1c.80.e5.04-
     .11.da.aa.ab.00.-
     15.58.0b.d9.47
VERSION: 6
RELEASE: 3
LEVEL: 1
VENDOR_D: IBM(R)
BRAND_D: POWER5(TM) QCM
TYPE_D: Quad-core Module
MODEL_D: All Existing
PRODUCT_D: IBM Spectrum Protect Extended Edition

```

Field descriptions

PRODUCT

Rollup of license types into products at the level presented in the QUERY PVUESTIMATE command. Possible values are PRODEE, PROTBASIC, PRODDATARET, PRODMAIL, PRODDDB, PRODSYSB, PRODSPACE, PRODSAN, PRODERP, or blank.

LICENSE_NAME

The license assigned to this node.

NODE_NAME

The node name.

LAST_USED

Date and time the identified node last connected to the system under this license.

TRYBUY

Indicates if running under try and buy mode. Possible values are TRUE or FALSE.

PROC_VENDOR

Processor vendor name as reported by the client.

PROC_BRAND

Processor brand name as reported by the client.

PROC_TYPE

Processor type as reported by the client. This value also reflects the number of cores. Example values are 1=SINGLE CORE, 2=DUO CORE, and 4=QUAD CORE.

PROC_MODEL
Processor model as reported by the client.

PROC_COUNT
Processor quantity.

ROLE
Node role. Possible values are CLIENT, SERVER, or OTHER.

ROLE_OVERRIDE
Override value specified in the UPDATE NODE command.

ROLE_EFFECTIVE
Actual role based on the values in the ROLE and ROLE_OVERRIDE fields.

VALUE_UNITS
Assigned processor value unit (PVU) for the processor.

PVU
Calculated PVU value.

$$\text{PVU per node} = \text{number of processors per node} * \text{processor type} * \text{pvu value}$$

where the `processor type` represents the number of cores, and the `pvu value` is the value defined for the processor type in the IBM® PVU table.

VALUE_FROM_TABLE
Flag that indicates whether the PVU was calculated based on the IBM PVU table. Possible values are YES or NO. If NO, a value of 100 is applied for each node defined as a server. If no role is defined for a node, the role of server is assumed for purposes of PVU calculation.

SCAN_ERROR
Flag that indicates whether license information was reported by client. Possible values are YES or NO.

API_CLIENT
Flag that indicates an API application. Possible values are YES or NO.

PVU_AGNOSTIC
Flag indicating that the client version release level is earlier than IBM Spectrum Protect V6.3. If the version is earlier than 6.3, valid PVU metrics are not expected. Possible values are YES or NO.

HYPERVISOR
Name of the virtual machine software as reported by the client.

GUID
Globally Unique Identifier (GUID) of the computer where the node is located. The GUID is obtained from the node table.

VERSION
Version of client.

RELEASE
Release of client.

LEVEL
Level of client.

VENDOR_D
Processor vendor display value from the PVU table.

BRAND_D
Processor brand display value from the PVU table.

TYPE_D
Processor type display value from the PVU table.

MODEL_D
Processor model display value from the PVU table.

PRODUCT_D
Product display value from the PVU table. The following values are possible:

- IBM Spectrum Protect
- IBM Spectrum Protect Extended Edition
- IBM Spectrum Protect for Data Retention
- IBM Spectrum Protect for SAN
- IBM Spectrum Protect for Space Management
- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Spectrum Protect for System Backup and Recovery
- Blank

Example: Obtain role and PVU-related information


The following example shows partial results for a selected node, including PVU-related information and role information. Possible roles are CLIENT, SERVER, or OTHER. PVU is calculated only for nodes defined as servers.

```
select * from nodes

ROLE: CLIENT
  ROLE_O: USERREPORTED
  PVENDOR: INTEL
  PBRAND: INTEL
  PTYPE: 4
  PMODEL:
  PCOUNT: 1
HYPERVISOR:
  PAPI: NO
  SCANNERERROR: NO
```

SET commands

Use the SET commands to specify values that affect many different IBM Spectrum Protect™ operations.

- SET ACCOUNTING (Set accounting records on or off)
- SET ACTLOGRETENTION (Set the retention period or the size of the activity log)
- SET ALERTACTIVEDURATION (Set the duration of an active alert)
- SET ALERTCLOSEDDURATION (Set the duration of a closed alert)
- SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)
- SET ALERTEMAILFROMADDR (Set the email address of the sender)
- SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)
- SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)
- SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)
- SET ALERTMONITOR (Set the alert monitor to on or off)
- SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)
- SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)
- SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)
- SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)
- SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)
- SET CLIENTACTDURATION (Set the duration period for the client action)
- SET CONFIGMANAGER (Specify a configuration manager)
- SET CONFIGREFRESH (Set managed server configuration refresh)
- SET CONTEXTMESSAGING (Set message context reporting on or off)
- SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)
- SET CROSSDEFINE (Specifies whether to cross-define servers)
- SET DBRECOVERY (Set the device class for automatic backups)
- SET DEDUPVERIFICATIONLEVEL (Set the percentage of extents to verify)
- SET DEFAULTAUTHENTICATION (Set the default authentication method for REGISTER NODE and REGISTER ADMIN commands)
- SET DEPLOYPKGMR (Enable the deployment package manager)
- SET DEPLOYREPOSITORY (Set the download path for client deployment packages)
- SET DEPLOYMAXPKGS (Set the maximum number of client deployment packages to store)
- SET DISSIMILARPOLICIES (Enable the policies on the target replication server to manage replicated data)
- SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)
- SET DRMCHECKLABEL (Specify label checking)
- SET DRMCMDFILENAME (Specify the name of a file to contain commands)
-  SET DRMCOPYCONTAINERSTGPOOL (Specify the container-copy storage pools to be processed by DRM commands)
- SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)
- SET DRMCOURIERNAME (Specify the courier name)
- SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)
- SET DRMFILEPROCESS (Specify file processing)
- SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)
- SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)
- SET DRMPLANPREFIX (Specify a prefix for recovery plan file names)

- SET DRMPLANVPOSTFIX (Specify replacement volume names)
- SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)
- SET DRMRPFEXPIREDAYS (Set criteria for recovery plan file expiration)
- SET DRMVaultNAME (Specify the vault name)
- SET EVENTRETENTION (Set the retention period for event records)
- SET FAILOVERHLADDRESS (Set a failover high level address)
- SET INVALIDPWLIMIT (Set the number of invalid logon attempts)
- SET LDAPPASSWORD (Set the LDAP password for the server)
- SET LDAPUSER (Specify an ID for an LDAP directory server)
- SET LICENSEAUDITPERIOD (Set license audit period)
- SET MAXCMDRETRIES (Set the maximum number of command retries)
- SET MAXSCHEDULESESSIONS (Set maximum scheduled sessions)
- SET MINPWLENGTH (Set minimum password length)
- SET MONITORINGADMIN (Set the name of the monitoring administrator)
- SET MONITOREDSEVERGROUP (Set the group of monitored servers)
- SET NODEATRISKINTERVAL (Specifies at-risk mode for an individual node)
- SET PASSEXP (Set password expiration date)
- SET PRODUCTOFFERING (Set the product offering that is licensed to your enterprise)
- SET QUERYSCHEDPERIOD (Set query period for polling client nodes)
- SET RANDOMIZE (Set randomization of scheduled start times)
- SET REPLRECOVERDAMAGED (Specify whether damaged files are recovered from a replication server)
- SET REPLRETENTION (Set the retention period for replication records)
- SET REPLSERVER (Set the target replication server)
- SET RETRYPERIOD (Set time between retry attempts)
- SET SCHEDMODES (Select a central scheduling mode)
- SET SERVERHLADDRESS (Set the high-level address of a server)
- SET SERVERLLADDRESS (Set the low-level address of a server)
- SET SERVERNAME (Specify the server name)
- SET SERVERPASSWORD (Set password for server)
- SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)
- SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)
- SET STATUSMONITOR (Specifies whether to enable status monitoring)
- SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)
- SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)
- SET SUBFILE (Set subfile backup for client nodes)
- SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)
- SET TAPEALERTMSG (Set tape alert messages on or off)
- SET TOCLOADRETENTION (Set load retention period for table of contents)
- SET VMATRISKINTERVAL (Specifies the at-risk mode for an individual VM filespace)

SET ACCOUNTING (Set accounting records on or off)

Use this command to determine whether an accounting record is created every time a client node session ends. An accounting record tracks the amount of storage used by a client node session.

Use the QUERY STATUS command to determine whether accounting records are generated. At installation, this value is set to OFF.

The accounting records are stored in an accounting file named dsmacct.log.

AIX | **Linux** The environment variable, DSMSERV_ACCOUNTING_DIR, specifies the directory where the accounting file is located.

Windows A registry entry controls the location of the accounting log.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ACCounting--+-ON-+-----><
```

Parameters

- ON
Specifies that the server creates an accounting record every time a client node session ends.
- OFF
Specifies that the server does not create accounting records.

Example: Create accounting records

To create an accounting record at the end of each client node session issue the command:

```
set accounting on
```

Related commands

Table 1. Commands related to SET ACCOUNTING

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET ACTLOGRETENTION (Set the retention period or the size of the activity log)

Use this command to manage the activity log records by date or size. The activity log contains normal activity messages generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

Activity log information includes messages, such as the following:

- Client session starts and ends
- Migration starts and ends
- Diagnostic error messages
- Scheduled administrative command output

At server installation, activity log management is retention-based, and the retention period is set to 30 days.

You can choose to adjust the length of time that the activity log retains messages to avoid insufficient or outdated data. The server automatically removes the messages from the activity log after the retention period passes.

Alternatively, you can choose to limit the total size of the activity log to control the amount of space occupied by the activity log. The server will periodically remove the oldest activity log records until the activity log size no longer exceeds the configured maximum size allowed.

You can issue the QUERY STATUS command to display the current number of records in the activity log and the size (in megabytes) that the activity log currently occupies.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>>Set ACTlogretention--number--+-Mgmtstyle-----Date-----+-----><
                                     '-Mgmtstyle-----+Date+-'
                                             '-Size-'

```

Parameters

number (Required)

Specifies the number of days to retain messages in the activity log when the log is managed by date, or specifies the maximum size of the activity log when it is managed by size. With retention-based management, a value of 1 specifies to retain the activity log records only for the current day. With size-based management, a value of 1 specifies a maximum size of 1 MB for the activity log. You can specify a number from 0 to 9999. A value of 0 disables activity log retention.

Mgmtstyle

Specifies whether activity log management is retention-based or size-based. This parameter is optional. The default is DATE. Possible values are:

Date

Specifies that activity log management is retention-based.

Size

Specifies that activity log management is size-based.

Example: Set the activity log retention period

Set the server to retain activity log records for 60 days. Issue the command:

```
set actlogretention 60
```

Example: Set the activity log size

Set the server to limit the size of the activity log to 300 MB. Issue the command:

```
set actlogretention 300 mgmtstyle=size
```

Related commands

Table 1. Command related to SET ACTLOGRETENTION

Command	Description
QUERY ACTLOG	Displays messages from the server activity log.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET ALERTACTIVEDURATION (Set the duration of an active alert)

Use this command to specify how long an alert remains active before it becomes inactive. If an active alert is triggered again, the duration is restarted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>--Set ALERTACTiveduration -number_mins-----<<
```

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains active before it becomes inactive. Specify a value from 1 to 20160. The initial server default value is 480 minutes.

Set the duration of an active alert to one day

Issue the following command to specify that alerts remain active for 1440 minutes before they change to inactive status:

```
set alertactiveduration 1440
```

Related commands

Table 1. Commands related to SET ALERTACTIVEDURATION

Command	Description
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	Specifies how long an alert remains inactive before it is closed.
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	Specifies how long an alert remains closed before it is deleted.
SET ALERTMONITOR (Set the alert monitor to on or off)	Specifies whether alert monitoring is set to on or off.
SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTCLOSEDDURATION (Set the duration of a closed alert)

Use this command to specify how long an alert remains closed before it is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTClosedduration -number_mins-----><
```

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains closed before it is deleted. Setting the value to 0 causes alerts to be deleted immediately after they are closed. Specify a value from 0 to 99999. The default value is set to 60 minutes when the IBM Spectrum Protect™ server database is initially formatted.

Delete alerts two hours after they are closed

Specify that alerts remain closed for 120 minutes before they are deleted:

```
set alertclosedduration 120
```

Related commands

Table 1. Commands related to SET ALERTCLOSEDDURATION

Command	Description
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET ALERTACTIVEDURATION (Set the duration of an active alert)	Specifies how long an alert remains active before it is moved to inactive status.
SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	Specifies how long an alert remains inactive before it is closed.
SET ALERTMONITOR (Set the alert monitor to on or off)	Specifies whether alert monitoring is set to on or off.

Command	Description
SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)

Use this command to enable alerts to be sent to specified administrators by email.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTEMail---ON--------><
      '-Off-'
```

Parameters

ON

Specifies that alerts can be sent to specified administrators by email.

OFF

Specifies that alerts cannot be sent to specified administrators by email. When the server database is initially formatted, the ALERTEMAIL setting is set to OFF.

Enable alerts to be sent to the administrator when they occur

Enable alerts to be sent by email by issuing the following command:

```
SET ALERTEMAIL ON
```

Related commands

Table 1. Commands related to SET ALERTEMAIL

Command	Description
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET ALERTEMAILFROMADDR (Set the email address of the sender)	Specifies the email address of the alert sender.
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	Specifies the SMTP mail server host name that is used to send alerts by email.
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	Specifies the SMTP mail server port that is used to send alerts by email.
SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILFROMADDR (Set the email address of the sender)

Use this command to specify the email address of the alert sender.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTEMAILFROMaddr -email_address-----<<
```

Parameters

email_address (Required)

Specifies the email address of the sender. Email addresses are in the form of *name@domain*. Email names, including the address, cannot exceed 64 characters in length, and the domain name cannot exceed 255 characters in length.

Specify the email address of the alert sender

Specify the email address of the sender by issuing the following command:

```
set alertemailfromaddr djadmin@mydomain.com
```

Related commands

Table 1. Commands related to SET ALERTEMAILFROMADDR

Command	Description
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	Enables alerts to be sent by email to specified administrators.
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	Specifies the SMTP mail server host name that is used to send alerts by email.
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	Specifies the SMTP mail server port that is used to send alerts by email.
SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)

Use this command to specify the Simple Mail Transfer Protocol (SMTP) mail server host name that is used to send the alert email.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTEMAILSMTPHost--host_name-----<<
```

Parameters

host_name (Required)

Specifies the SMTP mail server host name.

Specify the host name for the SMTP mail server as mail.domain.com

Specify *mail.domain.com* as the SMTP mail server, by issuing the following command:

```
set alertemailsmtp host mail.domain.com
```

Related commands

Table 1. Commands related to SET ALERTEMAILSMTPHOST

Command	Description
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	Enables alerts to be sent by email to specified administrators.
SET ALERTEMAILFROMADDR (Set the email address of the sender)	Specifies the email address of the alert sender.
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	Specifies the SMTP mail server port that is used to send alerts by email.
SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)

Use this command to specify the port number for the SMTP mail server. This mail server is used to send the alerts by email.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTEMAILSMTPPort--tcp_port-----<<
```

Parameters

tcp_port (Required)

Specifies the port number of the SMTP mail server. Specify a value of 1 through 32767. The default port number is 25.

Specify the port number of the SMTP mail server

Specify port number 450 as your SMTP mail server by issuing the following command:

```
set alertemailsmtpport 450
```

Related commands

Table 1. Commands related to SET ALERTEMAILSMTPPORT

Command	Description
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	Enables alerts to be sent by email to specified administrators.
SET ALERTEMAILFROMADDR (Set the email address of the sender)	Specifies the email address of the alert sender.
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	Specifies the SMTP mail server host name that is used to send alerts by email.
SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	Specifies the administrators that want to receive alert summaries by email.

SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)

Use this command to specify the administrators that want to receive alert summaries by email, every hour.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTSUMMARYToadmins---+admin_name+-----><
      ',-----'
```

Parameters

admin_name (Required)

Specifies the administrator name that wants to receive alert summaries by email. You can specify up to three administrator names by separating them with commas and no intervening spaces.

Specify two administrators to receive alert summaries

Specify that administrators HARRY and COLIN want to receive alert summaries, by issuing the following command:

```
set alertsummarytoadmins HARRY,COLIN
```

Related commands

Table 1. Commands related to SET ALERTSUMMARYTOADMINS

Command	Description
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	Enables alerts to be sent by email to specified administrators.
SET ALERTEMAILFROMADDR (Set the email address of the sender)	Specifies the email address of the alert sender.
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	Specifies the SMTP mail server host name that is used to send alerts by email.
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	Specifies the SMTP mail server port that is used to send alerts by email.

SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)

Use this command to specify how long an alert remains inactive. After the inactive duration is past, the alert is closed.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTINactiveduration -number_mins-----><
```

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains inactive before it is closed. You can specify a value in the range 1 - 20160. The initial server default value is 480 minutes.

Change alert status from inactive to closed after 60 minutes

Issue the following command to specify that an alert remains in inactive status for 60 minutes before it changes to closed status:

```
set alertinactiveduration 60
```

Related commands

Table 1. Commands related to SET ALERTINACTIVEDURATION

Command	Description
SET ALERTACTIVEDURATION (Set the duration of an active alert)	Specifies how long an alert remains active before it is moved to inactive status.
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	Specifies how long an alert remains closed before it is deleted.
SET ALERTMONITOR (Set the alert monitor to on or off)	Specifies whether alert monitoring is set to on or off.
SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTMONITOR (Set the alert monitor to on or off)

Use this command to turn the alert monitor on or off.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
                .-OFF-.  
>>-Set ALERTMONITOR -+-ON--+-+-----<<
```

Parameters

ON

Specifies that the IBM Spectrum Protect™ server monitors alerts.

OFF

Specifies that the IBM Spectrum Protect server does not monitor alerts. When the IBM Spectrum Protect server database is initially formatted, the alert monitoring setting is set to OFF.

Turn on alert monitoring

Turn on alert monitoring by issuing the following command:

```
set alertmonitor on
```

Related commands

Table 1. Commands related to SET ALERTMONITOR

Command	Description
SET ALERTACTIVEDURATION (Set the duration of an active alert)	Specifies how long an alert remains inactive before it is closed.
SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	Specifies how long an alert remains inactive before it is closed.
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	Specifies how long an alert remains closed before it is deleted.
SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)

Use this command to specify how often the alert monitor updates and prunes alerts that are stored in the IBM Spectrum Protect™ server database.

During this check interval, the alert monitor examines each alert on the server and completes the following actions:

- The alert monitor determines whether the active or inactive durations elapsed. If the specified duration elapses, the alert status is updated to the next state. For example:
 - Active to Inactive
 - Inactive to Closed
- If an alert is closed for the duration that is specified by the SET ALERTCLOSEDDURATION command, the alert is deleted.

You can use the QUERY MONITORSETTINGS command to determine whether alert monitoring is on. Use the SET ALERTMONITOR command to turn on alert monitoring.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ALERTUPDateinterval -number_mins-----<<
```

Parameters

number_mins (Required)

Specifies the length of time, in minutes, that the monitor waits before alerts are updated and pruned on the server. Specify a value from 1 to 9999. The server has an initial default value of 10 minutes.

Set alert update interval to 60 minutes

Specify that alerts are updated every hour by issuing the following command:

```
set alertupdateinterval 60
```

Related commands

Table 1. Commands related to SET ALERTUPDATEINTERVAL

Command	Description
SET ALERTACTIVEDURATION (Set the duration of an active alert)	Specifies how long an alert remains active before it is moved to inactive status.
SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	Specifies how long an alert remains inactive before it is closed.
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	Specifies how long an alert remains closed before it is deleted.
SET ALERTMONITOR (Set the alert monitor to on or off)	Specifies whether alert monitoring is set to on or off.

SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)

Use this command to activate and deactivate archive data retention protection. The server cannot contain any data in order for this command to work. At installation, the value is set to OFF.

When archive data retention protection is active:

- Only archive copies can be stored on the server.
- No archive copy can be deleted until the RETVER parameter in the DEFINE COPYGROUP (archive) command is satisfied.

Defining storage pools of type RECLAMATIONTYPE=SNAPLOCK is only supported on servers with data retention protection enabled.

Use the QUERY STATUS command to display the status of archive data retention protection.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-Set ARCHIVERETENTIONPROTECTION -+-OFF+-----><
                                     '-ON--'
```

Parameters

- OFF
Specifies that archive data retention protection is not active.
- ON
Specifies the archive data retention protection is active.

Example: Activate data retention protection

Activate archive data retention protection by issuing the following command:

```
set archiveretentionprotection on
```

Related commands

Table 1. Commands related to SET ARCHIVERETENTIONPROTECTION

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
AUDIT VOLUME	Compares database and storage pool information, and optionally, resolves any inconsistencies.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)

Use this command to set the server replication rule for archive data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for archive data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.

You can specify a normal-priority replication rule or a high-priority replication rule. In a replication process that includes both normal-priority and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain archive data and backup data. Replication of the archive data is a higher priority than the backup data. To prioritize the archive data, issue the SET ARREPLRULEDEFAULT command and specify the ALL_DATA_HIGH_PRIORITY replication rule. To prioritize the backup data, issue the SET BKREPLRULEDEFAULT command and specify the ALL_DATA replication rule for backup data. The ALL_DATA rule for backup data replicates backup data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set ARREPLRuledefault---ALL_DATA-----+-----><
++ALL_DATA_HIGH_PRIORITY--+
'-NONE-----'
```

Parameters

ALL_DATA
Replicates archive data with a normal priority.

ALL_DATA_HIGH_PRIORITY
Replicates archive data with a high priority.

NONE
Archive data is not replicated.

Example: Set the server replication rule for archive data

Set up the default rule for archive data to replicate with a high priority.

```
set arreplruledefault all_data_high_priority
```

Related commands

Table 1. Commands related to SET ARREPLRULEDEFAULT

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.

Command	Description
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)

Use this command to set the server replication rule for backup data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for backup data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.

You can specify normal-priority replication rules or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain archive data and active backup data. Replication of the active backup data is a higher priority than the archive data. To prioritize the backup data, issue the SET BKREPLRULEDEFAULT command and specify the ACTIVE_DATA_HIGH_PRIORITY replication rule. To prioritize the archive data, issue the SET ARREPLRULEDEFAULT command and specify the ALL_DATA replication rule for archive data. The ALL_DATA rule for archive data replicates archive data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set BKREPLRuledefault---ALL_DATA-----+----->>
      +-ACTIVE_DATA-----+
      +-ALL_DATA_HIGH_PRIORITY----+
      +-ACTIVE_DATA_HIGH_PRIORITY--+
      '-NONE-----'
```

Parameters

ALL_DATA

Replicates active and inactive backup data. The data is replicated with normal priority.

ACTIVE_DATA

Replicates active backup data. The data is replicated with normal priority.

Attention: If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the FORCERECONCILE=YES parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data. Data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority.

NONE

Backup data is not replicated.

Example: Set the server replication rule for backup data

Set up the default rule for backup data to replicate only active data and to replicate the data with a high priority.

```
set bkreplruledefault active_data_high_priority
```

Related commands

Table 1. Commands related to SET BKREPLRULEDEFAULT

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET REPLETENTION	Specifies the retention period for replication history records.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

SET CLIENTACTDURATION (Set the duration period for the client action)

Use this command to specify the duration for the schedule that was defined with the DEFINE CLIENTACTION command. A client action defines a schedule that runs one time on a client.

The program deletes these event records whether or not the client has processed the schedule. However, the schedules are not deleted until after the first event records are deleted. The retention period for events defaults to 10 days at installation.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET CLIENTACTDuration--days-----<<
```

Parameters

days (Required)

Specifies the number of days during which the schedule for the client action is active. You can specify an integer from 0 to 999. The default is 5 days.

The number of days you specify determines how long the database retains the schedule before deletion. A value of 0 indicates that the schedule duration is indefinite, and the schedule and associations are not deleted from the database.

Example: Set a 15-day duration period for the client action

To specify that the schedule for the client action be active for 15 days issue the following command.

```
set clientactduration 15
```

Related commands

Table 1. Commands related to SET CLIENTACTDURATION

Command	Description
DEFINE CLIENTACTION	Defines a command to be performed at a client node.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET CONFIGMANAGER (Specify a configuration manager)

Use this command to specify whether a server is a configuration manager. On a configuration manager, you can define configuration profiles to which other servers can subscribe.

You cannot designate a server as a configuration manager if the server subscribes to one or more profiles on another configuration manager.

If a server is a configuration manager, you cannot change this designation until you delete all profiles, including the default profile.

Issue the QUERY STATUS command to determine if a server is a configuration manager. When a server is installed, it is not designated as a configuration manager.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set CONFIGManager--+-OFF-+----->>  
                        '-ON--'
```

Parameters

ON

Specifies that the server is a configuration manager.

When you designate a server as a configuration manager, IBM Spectrum Protect™ creates a default profile named DEFAULT_PROFILE and associates with the profile all servers and server groups defined on the configuration manager. You can modify or delete the default profile.

OFF

Specifies that the server is not a configuration manager.

Example: Specify a configuration manager

Designate a server as a configuration manager.

```
set configmanager on
```

Related commands

Table 1. Commands related to SET CONFIGMANAGER

Command	Description
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET CONFIGREFRESH	Specifies a time interval for managed servers to contact configuration managers.

SET CONFIGREFRESH (Set managed server configuration refresh)

Use this command on a managed server to specify how often that server contacts its configuration manager for updated configuration information.

To display the current setting, issue the QUERY STATUS command. At installation, the interval is set to 60 minutes.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set CONFIGRefresh--minutes-----<<
```

Parameters

minutes (Required)

Specifies the interval, in minutes, before a managed server contacts its configuration manager for configuration updates. Specify an integer from 0 to 10000.

- If the value is greater than 0, the managed server immediately contacts the configuration manager. The next contact occurs when the specified interval is reached.
- If the value is 0, the managed server does not contact the configuration manager.

This value is ignored if the server does not subscribe to at least one profile on a configuration manager.

Example: Set a 45-minute refresh interval

Specify that a managed server contacts its configuration manager every 45 minutes.

```
set configrefresh 45
```

Related commands

Table 1. Commands related to SET CONFIGREFRESH

Command	Description
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UPDATE PROFILE	Changes the description of a profile.

SET CONTEXTMESSAGING (Set message context reporting on or off)

Use this command to get additional information when ANR9999D messages occur. IBM Spectrum Protect™ polls the server components for information that includes process name, thread name, session ID, transaction data, locks that are held, and database tables that are in use.

Note: When consecutive messages are issued from the same code area by the same thread, only the first of these messages will report the context information.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set CONTEXTmessaging--+-ON--+-----><
                               '-OFF-'
```

Parameters

- ON
Specifies to enable message context reporting.
- OFF
Specifies to disable message context reporting.

Example: Set message context reporting on or off

Turn on context messaging to receive additional information that could help determine the cause of ANR9999D messages.

```
set contextmessaging on
```

Related commands

Table 1. Commands related to SET CONTEXTMESSAGING

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)

Use this command to specify the number of days between client scans of workstation information that is used to estimate the processor value unit (PVU).

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set CPUINFOREFRESH--days-----><
```

Parameters

days (Required)

Specifies the number of days between scans for client devices. To retrieve the current setting, issue the QUERY STATUS command. The possible values are 1 - 9999. The default is 180.

Example: Set the amount of time before the next refresh to 90 days

```
SET CPUINFOREFRESH 90
```

Related commands

Table 1. Commands related to SET CPUINFOREFRESH

Command	Description
QUERY PVUESTIMATE	Displays an estimate of the client-devices and server-devices being managed.

SET CROSSDEFINE (Specifies whether to cross-define servers)

Use this command to specify whether a server is automatically defined to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set CROSSDefine---ON---+----->>  
          '-OFF-'
```

Parameters

ON

Specifies that a server may be cross-defined to another server. To automatically define one server to another, you must also permit cross defining in the server definition.

OFF

Specifies that a server may not be cross-defined to another server.

Example: Specifies whether to cross-define servers

Set cross define on to allow a server to be cross-defined to another server.

```
set crossdefine on
```

Related commands

Table 1. Command related to SET CROSSDEFINE

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERPASSWORD	Specifies the server password.

SET DBRECOVERY (Set the device class for automatic backups)

Use this command to specify the device class and number of data streams to be used for automatic database backups. You can also use this command to configure the BACKUP DB command to automatically back up the master encryption key for the server.

The master encryption key is used to encrypt data in directory-container and cloud-container storage pools, and to encrypt sensitive information in the server database. If you do not back up the master encryption key, you might not be able to access any of these encrypted items if a disaster occurs.

If you run the BACKUP DB command, and the device class is not the one that is specified in the SET DBRECOVERY command, a warning message is returned. However, the backup operation continues and is not affected.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
>>-SET DBRECOVery--device_class_name----->
      .-NUMStreams----1----- . -COMPRESS----No----- .
>--+-----+-----+-----+----->
      '-NUMStreams----number-' '-COMPRESS----+No--+-'
                                     '-Yes-'

      .-PROTECTKeys----Yes----- .
>--+-----+-----+-----+----->
      '-PROTECTKeys----+No--+-'
                                     '-Yes-'

>--+-----+-----+-----+----->>
      '-PASSWORD---password_name-'
```

Parameters

device_class_name **(Required)**

Specifies the device class to use for database backups.

NUMStreams

Specifies the number of parallel data movement streams to use when you back up the database. The default value is 1, and the maximum number is 32. Increasing this value causes a corresponding increase in the number of database backup sessions to be used and in the number of drives to be used for the device class. A NUMSTREAMS value that is specified in the BACKUP DB command overrides any value set in the SET DBRECOVERY command. The NUMSTREAMS value is used for all types of database backups.

If a value is specified that is greater than the number of drives available for the device class, the number of available drives are used. The available drives are defined to the device class by the MOUNTLIMIT parameter or by the number of online drives for the specified device class. The session is displayed in the QUERY SESSION output.

If you increase the number of streams, more volumes are used from the corresponding device class for this operation.

Using more volumes might improve the speed of the database backups, but at the cost of more volumes that are not fully used.

COMPRESS

Specifies whether volumes are compressed during database backup processing. This parameter is optional. The default value is No. You can specify one of the following values:

No

Specifies that the volumes that are created by the BACKUP DB command are not compressed.

Yes

Specifies that the volumes that are created by the BACKUP DB command are compressed.

If you specify the COMPRESS parameter on the BACKUP DB command, it overrides any value that is set in the SET DBRECOVERY command. Otherwise, the value that is set in the SET DBRECOVERY command is used.

Restrictions:

- Use caution when you specify the COMPRESS parameter. Using compression during database backups can reduce the size of the backup files. However, compression can increase the time to complete database backup processing.
- Do not back up compressed data to tape. If your system environment stores database backups on tape, set the COMPRESS parameter to No in the SET DBRECOVERY and BACKUP DB commands.

PROTECTKeys

Specifies that database backups include a copy of the master encryption key for the server that is used to encrypt node passwords, administrator passwords, and storage pool data. The master encryption key is stored in the dsmkeydb files. If you lose the dsmkeydb files, nodes and administrators are unable to authenticate with the server because the server is unable to read the passwords that are encrypted by using the master encryption key. In addition, any data that is stored in an encrypted storage pool cannot be retrieved without the master encryption key. This parameter is optional. The default value is Yes. You can specify one of the following values:

No

Specifies that database backups do not include a copy of the master encryption key for the server.

Attention: If you specify PROTECTKEYS=NO, you must manually back up the master encryption key for the server and make the key available when you implement disaster recovery. You cannot recover from a disaster without the master encryption key.

Yes

Specifies that database backups include a copy of the master encryption key for the server.

Attention: If you specify PROTECTKEYS=YES, you must also specify the PASSWORD parameter.

PASSword

Specifies the password that is used to protect the database backups. By default, database backup operations are protected by using a password. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

Important: Ensure that you remember this password. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database.

Example: Specify a device class for database backups

Specify the DBBACK device class for database backups. Run the following command:

```
set dbrecovery dback
```

Example: Specify a device class and number of streams for database backups

Specify the DBBACK device class for database backups, and specify that the backup is to use two data movement streams. Run the following command:

```
set dbrecovery dback numstreams=2
```

AIX

Linux

Windows

Example: Protect storage pool encryption keys in database backups

Encrypt storage pool data by specifying that database backups include a copy of the master encryption key for the server. Run the following command:

```
set dbrecovery dback protectkeys=yes password=password_name
```

Related commands

Table 1. Commands related to SET DBRECOVERY

Command	Description
BACKUP DB	Backs up the IBM Spectrum Protect database to sequential access volumes.
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

SET DEDUPVERIFICATIONLEVEL (Set the percentage of extents to verify)

Use this command to verify extents sent to the server during client-side data deduplication.

A rogue application that resides on a client system and that imitates the client, API, or GUI application can initiate an attack on the server. To reduce server vulnerability to such attacks, you can specify a percentage of client extents for the server to verify.

If the server detects that a security attack is in progress, the current session is canceled. In addition, the setting of the DEDUPLICATION parameter on the REGISTER NODE command is changed. The setting is changed from CLIENTORSERVER to SERVERONLY. The SERVERONLY setting disables client-side data deduplication for that node.

The server also issues a message that a potential security attack was detected and that client-side data deduplication was disabled for the node. If client-side data deduplication is disabled, all other client operations (for example, backup operations) continue. Only client-side data deduplication is disabled. If client-side data deduplication is disabled for a node because a potential attack was detected, the server deduplicates the data that is eligible for client-side data deduplication.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DEDUPVERificationlevel--+-0-----+-----><
                              '-percent_value-'
```

Parameters

percent_value (Required)

Specify an integer value 0 - 100 to indicate the percentage of client extents to be verified. A value of 0 indicates that no client extents are verified. The default for this command is 0.

Tips:

- Verifying extents consumes processing power and adversely affects server performance. For optimal performance, do not specify values greater than 10 for this command.
- To display the current value for SET DEDUPVERIFICATIONLEVEL, issue the QUERY STATUS command.

Example: Specify a minimum level of data deduplication verification

To specify that 1% of extents created during client-side data deduplication are verified, issue the following command:

```
set dedupverificationlevel 1
```

Example: Turn off data deduplication verification

To specify that none of the extents created during client-side data deduplication are verified, issue the following command:

```
set dedupverificationlevel 0
```

Related commands

Table 1. Commands related to SET DEDUPVERIFICATIONLEVEL

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE NODE	Changes the attributes that are associated with a client node.
UPDATE STGPOOL	Changes the attributes of a storage pool.

SET DEFAULTAUTHENTICATION (Set the default authentication method for REGISTER NODE and REGISTER ADMIN commands)

Use this command to set the default password authentication method for nodes and administrators that are the result of REGISTER NODE or REGISTER ADMIN commands.

If you specify LDAP, you establish the default value for authenticating to an external directory for any new REGISTER NODE or REGISTER ADMIN commands. This command makes it easier to register nodes or administrators when you use an LDAP directory server.

Tip: The default authentication setting can be overwritten when the authentication method is specified in a REGISTER NODE or REGISTER ADMIN command.

Privilege class

To issue this command you must have system privilege.

Syntax

```
>>-SET DEFAULTAUTHentication---Local+-----><
                    '-LDap--'
```

Parameters

Local

Specifies that any future REGISTER NODE or REGISTER ADMIN commands that you issue use LOCAL as the default authentication parameter value. Locally-authenticated passwords are those stored on the IBM Spectrum Protect™ server. The passwords authenticated locally are not case sensitive.

LDap

Specifies that any future REGISTER NODE or REGISTER ADMIN commands that you issue use LDAP as the default authentication parameter value. LDAP-authenticated passwords are those stored on an LDAP directory server and are case sensitive.

Example: Set the default password authentication value to LDAP

Specify that any REGISTER NODE or REGISTER ADMIN commands that you issue authenticate passwords with an LDAP directory server.

```
set defaultauthentication ldap
```

Related commands

Table 1. Commands related to SET DEFAULTAUTHENTICATION

Command	Description
SET LDAPPASSWORD	Sets the password for the LDAPUSER.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REGISTER NODE	Defines a client node to the server and sets options for that user.

SET DEPLOYPKGMR (Enable the deployment package manager)

Use this command to enable or disable the deployment package manager. This component downloads client deployment packages from the FTP site for automatic installation by using the Operations Center.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET DEPLOYPKGMgr--ON--<
>>-SET DEPLOYPKGMgr--OFF--<
```

Parameters

ON

Specifies that the deployment package manager queries the FTP site for new deployment packages and downloads new packages as they become available. This is the default.

OFF

Specifies that the deployment package manager does not query the FTP site or download new packages. If you disable the deployment package manager while packages are downloading, the active download processes continue to run until they are completed.

Example: Disable the deployment package manager

Disable the deployment package manager by issuing the following command:

```
set deploypkgmgr off
```

Related commands

Table 1. Commands related to SET DEPLOYPKGMGR

Command	Description
QUERY MONITORSETTINGS	Displays information about monitoring alerts and server status settings.
SET DEPLOYREPOSITORY	Specifies the location where client deployment packages are downloaded.

SET DEPLOYREPOSITORY (Set the download path for client deployment packages)

Use this command to specify the location where the automated deployment process downloads the latest client deployment packages. The deployment packages are used to install updates on client systems.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET DEPLOYREpository--path_name-----<
```

Parameters

path_name (Required)

Specifies the fully qualified path name where deployment packages are downloaded. This path also specifies the location where the server places the files that represent the storage volumes for the client deployment device class. You must specify a path name. If you do not, the server does not download the deployment packages.

When you modify the location where update packages are stored, previously downloaded packages are deleted automatically. Server volumes are deleted as data is pruned or expired.

Important: Do not manually delete files with a file name extension of .BFS. BFS files are volumes that are managed by the server, and they contain archive data that is expired or pruned automatically.

Example: Specify a path name

Specify `/source/packages/` as the location where deployment packages are downloaded. The same location is used for the IBM_DEPLOY_CLIENT_IMPORT device class, which is used for client deployment.

```
set deployrepository /source/packages/
```

Related commands

Table 1. Commands related to SET DEPLOYREPOSITORY

Command	Description
QUERY MONITORSETTINGS	Displays information about monitoring alerts and server status settings.
SET DEPLOYMAXPKGS	Specifies the maximum number of client deployment packages that are downloaded and stored on the server.

SET DEPLOYMAXPKGS (Set the maximum number of client deployment packages to store)

Use this command to specify the maximum number of client installable deployment packages that are downloaded and stored on the server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET DEPLOYMAXPkgs--number-----><
```

Parameters

number

Specifies the maximum number of deployment packages that are stored in the deployment repository for each product version. The minimum number of packages is 1, and the maximum number is 4. If you decrease the number, older versions of the packages are removed the next time packages are refreshed. It can take up to one day for packages to refresh. The default number is 4.

Example: Specify the maximum number of deployment packages

Specify 3 as the maximum number of deployment packages that are downloaded and stored.

```
set deploymaxpkgs 3
```

Related commands

Table 1. Commands related to SET DEPLOYMAXPKGS

Command	Description
QUERY MONITORSETTINGS	Displays information about monitoring alerts and server status settings.
SET DEPLOYREPOSITORY (Set the download path for client deployment packages)	Specifies the location where client deployment packages are downloaded.

SET DISSIMILARPOLICIES (Enable the policies on the target replication server to manage replicated data)

Use the SET DISSIMILARPOLICIES command to enable the policies that are defined on the target replication server to manage replicated client-node data. If you do not use the policies on the target replication server, replicated client-node data is managed by policies on the source replication server.

Ensure that IBM Spectrum Protect™, Version 7.1.1 or later, is installed on the source and target replication servers before you issue this command. Issue this command on the source replication server.

Before you use the policies that are defined on a target replication server, you must issue the VALIDATE REPLPOLICY command for that target replication server. This command displays the differences between the policies for the client nodes on the source replication server and policies on the target replication server. You can modify the policies on the target replication server before you enable these policies to manage replicated client-node data.

To obtain the name of the target replication server for which you want to manage data and to check whether the policies on the target replication server are set to ON, use the QUERY REPLSERVER command. At installation, the value is set to OFF.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DISSIMILARPolicies--target_server_name--+-OFF-+-----><
                                     +-OFF-+
                                     '-ON--'
```

Parameters

target_server_name (Required)

Specifies the name of the target replication server for which you want to enable the policies.

ON

Specifies that replicated client-node data is managed by the policies that are defined on the target replication server.

OFF

Specifies that replicated client-node data is managed by the policies that are defined on the source replication server. Off is the default value.

Example: Use the policies on a target replication server

To managed replicated client-node data from the target replication server, CVTCVS_LXS_SRV2, issue the following command on the source replication server:

```
set dissimilarpolicies CVTCVS_LXS_SRV2 on
```

Related commands

Table 1. Commands related to SET DISSIMILARPOLICIES

Command	Description
QUERY REPLSERVER	Displays information about replicating servers.

Command	Description
VALIDATE REPLPOLICY	Verifies the policies on the target replication server.

SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)

Use this command to specify names of the active-data pools to be recovered after a disaster. IBM Spectrum Protect™ uses these names if the PREPARE , MOVE DRMEDIA, or QUERY DRMEDIA command does not include the ACTIVEDATASTGPOOL parameter.

By default, volumes in active-data pools are not eligible for processing by disaster recovery manager. To process active-data pool volumes, you must issue the SET DRMACTIVEDATASTGPOOL command, or you must use the ACTIVEDATASTGPOOL command-line parameter on the MOVE DRMEDIA, QUERY DRMEDIA, or PREPARE command.

Use the QUERY DRMSTATUS command to display the current settings.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-'.-----'.
      v          |
>>-Set DRMACTIVEDatastgpool----active-data_pool_name+-----><

```

Parameters

active-data_pool_name (Required)

Specifies the active-data pool names. Separate multiple names with commas with no intervening spaces. You can use wildcard characters. The specified names will overwrite any previous settings. If you enter a null string (""), all current names are removed, and no active-data pool volumes in MOUNTABLE state are processed if they were not explicitly entered as MOVE DRMEDIA , QUERY DRMEDIA, or PREPARE command parameters.

Example: Set an eligible active-data pool

Set ACTIVEDATAPOOL1 as the eligible active-data pool.

```
set drmactivedatapool activedatastgpool1
```

Related commands

Table 1. Commands related to SET DRMACTIVEDATASTGPOOL

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.

SET DRMCHECKLABEL (Specify label checking)

Use this command to specify whether IBM Spectrum Protect™ reads the labels of sequential media checked out by the MOVE DRMEDIA command. At installation, the value of the DRMCHECKLABEL is set to YES.

Use the QUERY DRMSTATUS command to check the current setting.

AIX | **Linux** This command does not apply to 349X device types.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMCHECKLabel--+-Yes-+-----><
                        +-Yes-+
                        '-No--'
```

Parameters

- Yes**
Specifies that IBM Spectrum Protect reads the labels of sequential media checked out by the MOVE DRMEDIA command.
- No**
Specifies that IBM Spectrum Protect does not read the labels of sequential media checked out by the MOVE DRMEDIA command.

Example: Specify no label checking

Specify that no label checking is completed.

```
set drmchecklabel no
```

Related commands

Table 1. Commands related to SET DRMCHECKLABEL

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMCMDFILENAME (Specify the name of a file to contain commands)

Use this command to name a file that can contain the commands created when the MOVE DRMEDIA or QUERY DRMEDIA commands are issued. If the SET DRMCMDFILENAME is not issued, the MOVE DRMEDIA or QUERY DRMEDIA command generates a file name.

Use the QUERY DRMSTATUS command to display the current command file name.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMCMDFilename--file_name-----><
```

Parameters

file_name (Required)

AIX | **Linux** Specifies a full path name for a file to contain the commands created by the MOVE DRMEDIA or QUERY DRMEDIA command.

Windows Specifies a full path name for a file to contain the commands created by the MOVE DRMEDIA or QUERY DRMEDIA command. The file name can be up to 259 characters.

Attention: If a file of the same name already exists, MOVE DRMEDIA or QUERY DRMEDIA command tries to use it, and the existing data is overwritten.

Example: Specify a file name to contain DRMEDIA commands

AIX | **Linux** Specify a file name of /adsm/drm/orm/exec.cmds.

```
set drmcmdfilename /adsm/drm/orm/exec.cmds
```

Windows Specify a file name of c:\drm\orm\exec.cmd.

```
set drmcmdfilename c:\drm\orm\exec.cmd
```

Related commands

Table 1. Commands related to SET DRMCMDFILENAME

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

AIX | **Linux** | **Windows**

SET DRMCOPYCONTAINERSTGPOOL (Specify the container-copy storage pools to be processed by DRM commands)

Use this command to specify the container-copy storage pools to be processed by the MOVE DRMEDIA or QUERY DRMEDIA command when that command does not include the COPYCONTAINERSTGPOOL parameter.

By default, volumes in container-copy storage pools are not processed by the MOVE DRMEDIA and QUERY DRMEDIA commands. To process the volumes, you must issue the SET DRMCOPYCONTAINERSTGPOOL command, or you must use the COPYCONTAINERSTGPOOL parameter on the MOVE DRMEDIA or QUERY DRMEDIA command.

Tip: To display the current settings, use the QUERY DRMSTATUS command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
          .-|-----  
          v |  
>>-Set DRMCOPYCONTAINERSTGPOOL---pool_name+-----><
```

Parameters

pool_name (Required)

Specifies the names of the container-copy storage pools. Separate multiple names with commas and no intervening spaces. You can use wildcard characters. The specified names replace any previous setting. If you enter a null string (""), all current names are removed.

Example: Specify storage pools to be processed by the MOVE DRMEDIA and QUERY DRMEDIA commands

Set CONTCOPY1 and CONTCOPY2 as the container-copy storage pools to be processed.

```
set drmcopystgpool contcopy1,contcopy2
```

Related commands

Table 1. Commands related to SET DRMCOPYCONTAINERSTGPOOL

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)

Use this command to specify names of the copy storage pools to be recovered after a disaster. IBM Spectrum Protect™ uses these names if the PREPARE command does not include the COPYSTGPOOL parameter.

If the MOVE DRMEDIA or QUERY DRMEDIA command does not include the COPYSTGPOOL parameter, the command processes the volumes in the MOUNTABLE state that are in the copy storage pool named by the SET DRMCOPYSTGPOOL command. At installation, all copy storage pools are eligible for DRM processing.

Use the QUERY DRMSTATUS command to display the current settings.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMCOPYstgpool-----copy_pool_name+-----<<
```

Parameters

copy_pool_name (Required)

Specifies the copy storage pool names. Separate multiple names with commas and no intervening spaces. You can use wildcard characters. The specified names replace any previous setting. If you enter a null string (""), all current names are removed, and all copy storage pools are eligible for processing.

Example: Set an eligible copy storage pool

Set COPYSTGPOOL1 as the eligible copy storage pool.

```
set drmcopystgpool copystgpool1
```

Related commands

Table 1. Commands related to SET DRMCOPYSTGPOOL

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.

SET DRMCOURIERNAME (Specify the courier name)

Use this command to specify the courier name. At installation, this name is set to COURIER. The MOVE DRMEDIA command uses the courier name to set the location of volumes that are moving to the COURIER state.

You can use the QUERY DRMSTATUS to see the name of the courier.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMCOURiername--courier_name-----><
```

Parameters

`courier_name` (Required)

Specifies the name of the courier. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Set the courier name

Set the name of the courier to Joe's Courier Service.

```
set drmcouriername "Joe's Courier Service"
```

Related commands

Table 1. Commands related to SET DRMCOURIERNAME

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)

Use this command to specify when a database backup series is eligible to be expired.

The value set by this command applies to both a snapshot and a full plus incremental database backup series. Any type of database backup series is eligible for expiration if all of the following are true:

- The age of the last volume of the series exceeds the expiration value set with the SET DRMDBBACKUPEXPIREDAYS command and the value that is specified for the DELgraceperiod parameter in the DEFINE SERVER command. The DELgraceperiod parameter applies only to remote database backups. The default value for the DELgraceperiod parameter is 5 days. For example, if you set the value for the SET DRMDBBACKUPEXPIREDAYS command to 7 days and set the value for the DELgraceperiod parameter to 6 days, the remote database backup series does not expire until 13 days elapse.
- For volumes that are not virtual volumes, all volumes in the series are in the VAULT state.
- The volume is not part of the most recent database backup series.

Remember: The most recent backup series of either type is not deleted.

See the MOVE DRMEDIA command for more information on the expiration of database backup volumes that are not virtual volumes. See the EXPIRE INVENTORY command for more information on expiration of database backup volumes that are virtual volumes.

Use the QUERY DRMSTATUS to see the number of days specified.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMDBBackupexpiredays--days-----<<
```

Parameters

days (Required)

Specifies the number of days that must elapse since a database series was created before it is eligible to be expired. The number of days must match the volume reuse delay period for copy storage pools that are managed by disaster recovery manager. Specify an integer value 0 - 9999.

Example: Set the database backup series expiration

Set the database backup series expiration value to 60.

```
set drmdbbackupexpiredays 60
```

Related commands

Table 1. Commands related to SET DRMDBBACKUPEXPIREDAYS

Command	Description
DSMSERV RESTORE DB	Restores an IBM Spectrum Protect database.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
DEFINE SERVER	Defines a server for server-to-server communications.

SET DRMFILEPROCESS (Specify file processing)

Use this command to specify if the MOVE DRMEDIA or QUERY DRMEDIA command should process database backup volumes and copy storage pool volumes that are associated with a FILE device class. At installation, the value is set to NO. Use the QUERY DRMSTATUS to determine the current setting.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMFILEProcess--+-No--+------<<
```

+-No--+
'-Yes-'

Parameters

No

Specifies that the MOVE DRMEDIA and QUERY DRMEDIA commands does not process database backup and copy storage pool volumes that are associated with a FILE device class. This is the default.

Yes

Specifies that the MOVE DRMEDIA and QUERY DRMEDIA commands process database backup and copy storage pool volumes that are associated with a FILE device class.

Example: Specify that the DRMEDIA commands do not include FILE type device classes

Set the file processing value to no.

```
set drmfileprocess no
```

Related commands

Table 1. Commands related to SET DRMFILEPROCESS

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)

Use this command to specify a prefix to the recovery instructions file name. If you issue this command, IBM Spectrum Protect™ uses the specified prefix if the PREPARE command is issued without the INSTRPREFIX parameter.

Use the QUERY DRMSTATUS command to display the current value for the prefix.

AIX | **Linux** the prefix is the current IBM Spectrum Protect server working directory.

Windows If no prefix is set, the prefix is set to the directory representing this instance of the server, which is typically the directory that the server was originally installed from.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMINSTRPrefix--prefix-----><
```

Parameters

AIX | **Linux** prefix (Required)
AIX | **Linux**

Specifies a path name prefix for the files that contain the recovery instructions. When processing the PREPARE command, IBM Spectrum Protect appends the name of the appropriate recovery plan file stanza to find the file. The maximum length is 250 characters.

The prefix can be one of the following:

- **Directory path:** End the prefix with a forward slash (/). For example:

```
/admsrv/recinstr/
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
/admsrv/recinstr/RECOVERY.INSTRUCTIONS.GENERAL
```

- **Directory path followed by a string:** IBM Spectrum Protect treats the string as part of the file name. For example:

```
/admsrv/recinstr/accounts
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
/admsrv/recinstr/accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

- **String only:** IBM Spectrum Protect specifies the directory path and appends the appropriate recovery plan file stanza name.
 - IBM Spectrum Protect uses the name of the current working directory. For example, the current working directory is /opt/tivoli/tsm/server/bin. You specify the following:

```
shipping
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would look like this:

```
/opt/tivoli/tsm/server/bin/shipping.RECOVERY.INSTRUCTIONS.GENERAL
```

Windows prefix (Required)

Windows

Specifies a path name prefix for the files that contain the recovery instructions. When processing the PREPARE command, IBM Spectrum Protect appends the name of the appropriate recovery plan file stanza to find the file. The maximum length is 200 characters.

The prefix can be one of the following:

- **Directory path:** End the prefix with a back slash (\). For example:

```
c:\admsrv\recinstr\
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
c:\admsrv\recinstr\RECOVERY.INSTRUCTIONS.GENERAL
```

- **Directory path followed by a string:** IBM Spectrum Protect treats the string as part of the file name. For example:

```
c:\admsrv\recinstr\accounts
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
c:\admsrv\recinstr\accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

- **String only:** IBM Spectrum Protect specifies the directory path and appends the appropriate recovery plan file stanza name. The directory path is the directory representing this instance of the IBM Spectrum Protect server (typically the original IBM Spectrum Protect server installation directory). For example, the directory representing this instance of the server is c:\Program Files\Tivoli\TSM;\server2, and you specify the following prefix:

```
shipping
```

The resulting recovery plan file name is:

```
c:\Program Files\Tivoli\TSM;\server2\shipping.19971115.051421
```

Example: Specify the recovery plan prefix

AIX | **Linux** Specify reading the recovery plan instructions from directory /drmpplan/primesrv.

```
set drminstrprefix /drmpplan/primesrv/
```

Windows Specify reading the recovery plan instructions from directory c:\win32app\ibm\adsm\server2\.

```
set drminstrprefix c:\win32app\ibm\adsm\server2\
```

Related commands

Table 1. Commands related to SET DRMINSTRPREFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)

Use this command to specify the name of the onsite location for storing the media. At installation, the name is set to NOTMOUNTABLE. Use the QUERY DRMSTATUS command to see the location name.

The location name is used by the MOVE DRMEDIA command to set the location of volumes that are moving to the NOTMOUNTABLE state.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMNOTMOUNTABLENAME--location-----<<
```

Parameters

location (Required)

Specifies the name of the onsite location for storing the media. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Specify the name of the onsite location

Set the name of the location to room 123/31.

```
set drmnotmountablename "room 123/31"
```

Related commands

Table 1. Commands related to SET DRMNOTMOUNTABLENAME

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPLANPREFIX (Specify a prefix for recovery plan file names)

Use this command to specify a prefix for a recovery plan file name.

If you issue this command, IBM Spectrum Protect™ uses the specified prefix if the PREPARE command does not include the PLANPREFIX parameter.

Use the QUERY DRMSTATUS command to display the current value for the recovery plan prefix.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMPLANPrefix--prefix----->>
```

Parameters

AIX | **Linux** prefix (Required)

AIX | **Linux** Specifies the prefix for a recovery plan file name. The maximum length of the prefix is 250 characters. If you enter a null string (""), the current prefix is removed, and the server uses the algorithm described in the PLANPREFIX parameter in the PREPARE command.

For the prefix, you can specify:

- **A directory path followed by a forward slash (/):** IBM Spectrum Protect appends to the prefix the date and time in the `yyyymmdd.hhmmss` format. For example, the SET DRMPLANPREFIX is set to the following:

```
/admsrv/recplans/
```

The resulting recovery plan file name is:

```
/admsrv/recplans/19971115.051421
```

- **A directory path followed by a string:** IBM Spectrum Protect uses the string as part of the file name. IBM Spectrum Protect appends to the prefix the date and time in the `.yyyymmdd.hhmmss` format (note the initial period). For example, the SET DRMPLANPREFIX is set to the following:

```
/admsrv/recplans/accounting
```

The resulting recovery plan filename is:

```
/admsrv/recplans/accounting.19971115.051421
```

- **A string that is not preceded by a directory path:** IBM Spectrum Protect appends to the prefix the date and time information in the `.yyyymmdd.hhmmss` format (note the initial period). IBM Spectrum Protect determines the directory path as follows:

- IBM Spectrum Protect uses the directory path name of the current working directory of the IBM Spectrum Protect server. For example, the current IBM Spectrum Protect working directory is `/opt/tivoli/tsm/server/bin`. The SET DRMPLANPREFIX command is set to the following:

```
shipping
```

The resulting recovery plan file name is:

```
/opt/tivoli/tsm/server/bin/shipping.19971115.051421
```

Windows prefix (Required)

Windows Specifies a prefix for the path name used to generate the recovery plan file name. The prefix can be up to 200 characters. IBM Spectrum Protect uses the prefix if the PREPARE command is issued without the PLANPREFIX parameter. IBM Spectrum Protect builds a unique recovery plan file name by appending to the prefix the date and time format: `yyyymmdd.hhmmss` (for example, 19951115.051421). If you enter a null string (""), the current prefix is removed, and the server uses the algorithm described in the PLANPREFIX parameter in the PREPARE command.

For the prefix, you can specify:

1. A directory path
2. A directory path followed by a string
3. A string

The following describes the rules for possible prefix specifications:

1. To specify a directory path for the prefix, end the prefix with a back slash (\). IBM Spectrum Protect appends to the prefix the date and time information using the `yyyymmdd.hhmmss` format. For example the SET DRMPLANPREFIX is set to the following:

```
c:\admsrv\recplans\
```

The resulting recovery plan file name is:

```
c:\admsrv\recplans\19951115.051421
```

Important: If you issue the SET DRMPLANPREFIX command from a command line client and the last character in the command line is a back slash, IBM Spectrum Protect interprets it as a continuation character. To avoid this, enclose the prefix in quotation marks. For example: "c:\admsrv\recplans\"

2. If the prefix is a directory path followed by a string, IBM Spectrum Protect uses the string as part of the file name. IBM Spectrum Protect appends to the prefix the date and time in the .yyyymmdd.hhmmss format (note the initial period). For example, the SET DRMPLANPREFIX is set to the following

```
c:\admsrv\recplans\accounting
```

The resulting recovery plan filename is the following:

```
c:\admsrv\recplans\accounting.19951115.051421
```

3. If the prefix is a string that is not preceded by a directory path, IBM Spectrum Protect appends to the prefix the date and time information in the .yyyymmdd.hhmmss format (note the initial period). The directory path that IBM Spectrum Protect uses is the directory path representing this instance of the IBM Spectrum Protect server (typically the directory that the IBM Spectrum Protect server was originally installed from). For example, the directory representing this instance of the server is c:\Program Files\Tivoli\TSM;\server2, and you set the prefix to:

```
shipping
```

The resulting recovery plan filename is:

```
c:\Program Files\Tivoli\TSM;\server2\shipping.19951115.051421
```

Example: Specify a prefix for recovery plan file names

Specify a prefix so that the generated recovery plan files are stored in the following directory:

- **AIX** **Linux** /drmpln/primsrv
- **Windows** c:\drmtest\prepare\

Issue the command: **AIX** **Linux**

```
set drmplnprefix /drmpln/primsrv/
```

Windows

```
set drmplnprefix c:\drmtest\prepare\
```

Related commands

Table 1. Commands related to SET DRMPLANPREFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPLANVPOSTFIX (Specify replacement volume names)

Use this command to specify the character to be appended to replacement volume names in the recovery plan file. The character can help you find or generate replacement volume names when you use the recovery plan file.

At installation, the character is set to @. IBM Spectrum Protect™ generates replacement names for primary storage pool volumes that were added by the DEFINE VOLUME command. Use the appended character to:

- Find replacement volume names in the recovery plan stanzas so that you can change the names at recovery time. For example, you may not know the names of the available tape volumes at the recovery site.
- Generate replacement volume names. You need a naming convention that works for any device type in your primary storage pools. Consider the following:
 - The generated length of replacement volume name
 - Legal characters in the replacement volume name

- Conflicts with existing volume names
- A replacement volume name must be different from any destroyed, existing, or new volume name.

Use the QUERY DRMSTATUS command to see the character added to the end of the replacement volume names.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMPLANVpostfix--character-----><
```

Parameters

character (Required)

Specifies the character appended to the replacement volume names in the recovery plan file. Specify an alphanumeric or special character.

AIX Attention: A special character can cause unpredictable results in the AIX® shell or command line environment.

Windows Attention: A special character can cause unpredictable results in the Windows batch/command line environment.

Example: Specify the appended character for replacement volume names

Set the character appended to the replace volume names to R.

```
set drmplnvpostfix R
```

Related commands

Table 1. Commands related to SET DRMPLANVPOSTFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)

Use this command to specify the names of primary storage pools that you want to recover. If the PREPARE command does not include the PRIMSTGPOOL parameter, DRM processes the names specified in this command.

Use the QUERY DRMSTATUS command to display the current settings. At installation, all primary storage pools defined to the server are eligible for DRM processing.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .,-----
      v          |
>>-Set DRMPRIMstgpool---primary_pool_name-+-----><

```

Parameters

primary_pool_name (Required)

Specifies the names of the primary storage pool names you want to recover. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. The names that you specify replace any previous setting. If you enter a null string (""), all current names are removed, and all primary storage pools are eligible for DRM processing.

Example: Set a primary storage pool to be managed by DRM

Set the primary storage pool to be managed by DRM to PRIMSTGPOOL1.

```
set drmprimstgpool primstgpool1
```

Related commands

Table 1. Commands related to SET DRMPRIMSTGPOOL

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.

SET DRMRPFEXPIREDAYS (Set criteria for recovery plan file expiration)

Use this command to specify when recovery plan files are eligible for expiration. This command and expiration processing apply only to recovery plan files that were created with the DEVCLASS parameter specified on the PREPARE command (that is, virtual volumes of type RPFIL and RPSNAPSHOT). Expiration processing on the source server expires plan files that are stored on the target server. Locally created recovery plan files are not expired.

An RPFIL file is associated with a full plus incremental database backup series. An RPSNAPSHOT file is associated with a database snapshot backup series.

Attention: The latest RPFIL and RPSNAPSHOT files are never deleted.

A recovery plan file is eligible for expiration if both of the following are true:

- The last recovery plan file of the series exceeds the expiration value that is specified with the SET DRMRPFEXPIREDAYS command and the value that is specified for the DELgraceperiod parameter in the DEFINE SERVER command. The default value for the DELgraceperiod parameter is 5 days. For example, if you set the value for the SET DRMRPFEXPIREDAYS command to 80 days and set the value for the DELgraceperiod parameter to 6 days, the recovery plan file does not expire until 86 days elapse.
- The latest recovery plan file is not associated with the most recent database backup series.

For more information about expiration processing, see the EXPIRE INVENTORY command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set DRMRPFExpiredays--days-----><
```

Parameters

days (Required)

Specifies the number of days that must elapse before a recovery plan file expires. You can specify a number 0 - 9999. At installation, this value is set to 60.

Example: Set the recovery plan expiration

Set the recovery plan file expiration value to 30.

Related commands

Table 1. Commands related to SET DRMRPFEXPIREDDAYS

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.
QUERY RPFFILE	Displays information about recovery plan files.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMDBBACKUPEXPIREDDAYS	Specifies criteria for database backup series expiration.
DEFINE SERVER	Defines a server for server-to-server communications.

SET DRMVaultName (Specify the vault name)

Use this command to specify the vault name. At installation the name is set to VAULT. Use the QUERY DRMSTATUS command to see the name of the vault.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET DRMVaultname--vault_name-----><
```

Parameters

vault_name (Required)

Specifies the name of the vault. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Specify a vault name

Specify ironmountain as the vault name.

```
set drmvaultname ironmountain
```

Related commands

Table 1. Commands related to SET DRMVaultName

Command	Description
MOVE DRMEDIA	Moves DRM media onsite and offsite.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET EVENTRETENTION (Set the retention period for event records)

Use this command to set the retention period for event records in the server database that will allow you to monitor completed schedules. An event record is created whenever processing of a scheduled command is started or missed.

You can adjust the length of time that the server maintains event information to avoid insufficient or outdated data. The server automatically removes the event records from the database after the retention period passes and the startup window for the event has elapsed.

You can issue the QUERY EVENT command to display information about scheduled and completed events.

You can issue the DELETE EVENT command to delete event records regardless of whether their retention period has passed.

You can issue the QUERY STATUS command to display the value for the event retention period. At installation, this value is set to 10 days.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set EVentretention--days-----><
```

Parameters

days (Required)

The number of days that the database retains event records. You can specify an integer from 0 to 9999. A value of 0 indicates that only event records for the current day are retained.

Example: Set the retention period for event records

Set the retention period to 15 days.

```
set eventretention 15
```

Related commands

Table 1. Commands related to SET EVENTRETENTION

Command	Description
DELETE EVENT	Deletes event records before a specified date and time.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET FAILOVERHLADDRESS (Set a failover high level address)

Use this command to specify the IP address that a client uses to connect to this server as the secondary replication server during failover, if the address is different from the IP address that is specified for the replication process.

You must specify the address of the server that is used if the high-level address (HLA) is different. This command is required only if you use separate dedicated networks for server-to-server communication and client access.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET FAILOVERHladdress--high_level_address-----><
```

Parameters

high_level_address (Required)

Specifies a server HLA as a numeric dotted decimal name or a host name to use during failover. If you specify a host name, a server that can resolve the name to the dotted decimal format must be available.

To remove the failover IP address, issue the command without specifying a value.

Example: Set a failover high-level address

The name of the HLA that you want to set for failover operations on this server.

```
set failoverhladdress server1
```

Example: Remove a high-level address

To remove a high-level address for a failover server, issue the following command:

```
set failoverhladdress
```

Related commands

Table 1. Commands related to QUERY REPLSERVER

Command	Description
QUERY REPLSERVER (Query a replication server)	Displays information about replicating servers.
REMOVE REPLSERVER (Remove a replication server)	Removes a server from replication.

SET INVALIDPWLIMIT (Set the number of invalid logon attempts)

Use this command to set the number of invalid logon attempts that are allowed before a node is locked.

The SET INVALIDPWLIMIT command also applies to LDAP directory servers that store complex node passwords. LDAP directory servers can limit the number of invalid password attempts independent of the IBM Spectrum Protect™ server. You might not want to set up the LDAP directory server for invalid attempts for the IBM Spectrum Protect namespace if you use the SET INVALIDPWLIMIT command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set--INVALIDPwlimit--number-----><
```

Parameters

number (Required)

Specifies the number of invalid logon attempts allowed before a node is locked.

You can specify an integer from 0 to 9999. A value of 0 means that invalid logon attempts are not checked. A value of 1 means that if a user issues an invalid password one time, the node is locked by the server. The default is 0.

Important: If your password is authenticated with an LDAP directory server, it can be managed by the LDAP server and the IBM Spectrum Protect server. Not all IBM Spectrum Protect server commands affect passwords that authenticate with an LDAP server. For example, the SET PASSEXP and RESET PASSEXP commands do not affect passwords that authenticate with an LDAP directory server. You can manage your password features through the IBM Spectrum Protect server. If you issued the SET INVALIDPWLIMIT command, all IBM Spectrum Protect passwords are controlled by the limit that you set. If you configure the LDAP directory server to limit the number of invalid password attempts, a conflict might occur.

Example: Define the number of allowed invalid login attempts

Set the number of invalid logon attempts allowed.

```
set invalidpwlimit 6
```

Related commands

Table 1. Commands related to SET INVALIDPWLIMIT

Command	Description
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET MINPWLENGTH	Sets the minimum length for client passwords.

SET LDAPPASSWORD (Set the LDAP password for the server)

Use this command to define a password for the user or account ID that you specified by using the SET LDAPUSER command.

Requirement: You must define the LDAPURL option and issue the SET LDAPUSER command before you issue the SET LDAPPASSWORD command. If the LDAPURL option is not defined when you set the user password for the Lightweight Directory Access Protocol (LDAP) server, you must restart the IBM Spectrum Protect™ server after you define the LDAPURL option.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set LDAPPASSWORD--ldap_user_password-----<<
```

Parameters

ldap_user_password

Specifies the password that the IBM Spectrum Protect server uses when it authenticates to the LDAP server. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters. If you have equal signs within your password, you must contain the whole password within quotation marks. You can use the following characters:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )  
| { } [ ] : ; < > , ? / ~
```

Example: Set an LDAP password

```
set ldappassword LdAp20&12PaSsWoRd
```

Example: Set an LDAP password that includes an equal sign

```
set ldappassword "LdAp=LastWoRd"
```

Related commands

Table 1. Commands related to SET LDAPPASSWORD

Command	Description
AUDIT LDAPDIRECTORY	Audit an IBM Spectrum Protect-controlled namespace on an LDAP directory server.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.

SET LDAPUSER (Specify an ID for an LDAP directory server)

Use this command to specify the ID of a user or account that can access a Lightweight Directory Access Protocol (LDAP) server.

The specified ID must have read access to the accounts on the LDAP server that are used for authentication. To modify LDAP IDs or reset passwords for LDAP IDs, the specified ID must have write authority for accounts on the LDAP server.

Tip: The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set LDAPUser--ldap_user_dn-----<<
```

Parameters

ldap_user_dn
Specifies the ID of a user or account that can access an LDAP server.

Example: Specify an administrative user ID for conducting operations on an LDAP server

To specify an administrator with a user ID of JACKSPRATT, who represents a US company that is named EXAMPLE, issue the following command:

```
set ldapuser JackSpratt@us.example.com
```

Related commands

Table 1. Commands related to SET LDAPUSER

Command	Description
AUDIT LDAPDIRECTORY	Audit an IBM Spectrum Protect-controlled namespace on an LDAP directory server.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.

SET LICENSEAUDITPERIOD (Set license audit period)

Use this command to specify the period, in days, between automatic license audits performed by IBM Spectrum Protect™.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set--LICenseauditperiod--+-30---+-----+----->>  
'-days-'
```

Parameters

days

Specifies the number of days between automatic server license audits. This parameter is optional. The default value is 30. You can specify an integer from 1 to 30, inclusive.

Example: Specify a 14 day server license audit

Specify that the server audits licenses every 14 days.

```
set licenseauditperiod 14
```

Related commands

Table 1. Commands related to SET LICENSEAUDITPERIOD

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY AUDITOCCUPANCY	Displays the server storage utilization for a client node.
QUERY LICENSE	Displays information about licenses and audits.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the IBM Spectrum Protect server.

SET MAXCMDRETRIES (Set the maximum number of command retries)

Use this command to set the maximum number of times that a scheduler on a client node can retry a failed, scheduled command.

You can use the command to override the maximum number of retries that are specified by the client node. A client's value is overridden only if the client is able to connect with the server.

This command is used with the SET RETRYPERIOD command to regulate the time and the number of retry attempts to rerun failed command.

You can issue the QUERY STATUS command to display the current retry value. At installation, IBM Spectrum Protect™ is configured so that each client determines its own retry value.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set MAXCMDRetries--+-number-+-----+----->>  
'-number-'
```

Parameters

number

Specifies the maximum number of times the scheduler on a client node can retry a failed scheduled command. This parameter is optional.

The default is that each client determines its own value for this parameter. You can specify an integer from 0 to 9999. See the appropriate client documentation for more information on setting the maximum command retries from the client.

Example: Set the maximum number of command retries to 2

Retry, only twice, a failed attempt to process a scheduled command.

```
set maxcmdretries 2
```

Related commands

Table 1. Command related to SET MAXCMDRETRIES

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.

SET MAXSCHEDESESSIONS (Set maximum scheduled sessions)

Use this command to set the number of sessions that the server can use to process scheduled operations. This command specifies the maximum number of scheduled sessions as a percentage of the total number of available server sessions.

Limiting the number of sessions ensures that some are available for unscheduled operations, such as backup or archive. You can increase either the total number of sessions (with the MAXSESSIONS parameter) or the maximum percentage of scheduled sessions. Increasing the total number of sessions available, however, can affect server performance. Increasing the maximum percentage of scheduled sessions can reduce the sessions available for unscheduled operations.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set MAXSCHeDsessions--percent-----<<
```

Parameters

percent (Required)

Specifies the percentage of total server sessions that can be used for scheduled operations. You can specify an integer from 0 to 100. The MAXSESSIONS parameter in the server options file determines the maximum number of total available server sessions.

If you set the maximum percentage of scheduled sessions to 0, no scheduled events can begin. If you set the maximum percentage of scheduled sessions to 100, the maximum number of scheduled sessions is the value of the MAXSESSIONS option.

Tip: If the maximum number of scheduled sessions do not coincide with the percentage that you set in the SET MAXSCHEDESESSIONS command, run the SET MAXSCHEDESESSIONS command again. Look in the MAXSESSIONS option and determine the number that is specified there. If the MAXSESSIONS option number changed and you did not issue the SET MAXSCHEDESESSIONS command since the change, the maximum number of scheduled sessions can change.

Set a maximum of 20 sessions for scheduled activities

The MAXSESSIONS option has a value of 80. If you want no more than 20 sessions to be available for scheduled activity, set the percentage to 25.

Related commands

Table 1. Commands related to SET MAXSCHEDESESSIONS

Command	Description
QUERY OPTION	Displays information about server options.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET MINPWLENGTH (Set minimum password length)

Use this command to set the minimum length of a password.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set--MINPwlength--+-8-----+----->>
                        '-length-'
```

Parameters

length (Required)

Specifies the minimum length of a password. This parameter is optional. You can specify an integer in the range 1 - 64. The default value is 8.

Example: Set the minimum password length

Set the minimum password length to 12 characters.

```
set minpwlenth 12
```

Related commands

Table 1. Commands related to SET MINPWLENGTH

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.

Related reference:

SET SERVERPASSWORD (Set password for server)
 DEFINE SERVER (Define a server for server-to-server communications)
 UPDATE SERVER (Update a server defined for server-to-server communications)
 REGISTER ADMIN (Register an administrator ID)
 UPDATE ADMIN (Update an administrator)
 REGISTER NODE (Register a node)
 UPDATE NODE (Update node attributes)
 SET LDAPPASSWORD (Set the LDAP password for the server)
 BACKUP DB (Back up the database)
 SET DBRECOVERY (Set the device class for automatic backups)

SET MONITOREDSEVERGROUP (Set the group of monitored servers)

Use this command to set the group of servers that are being monitored for alerts and status. You can also use this command to change or remove the group of monitored servers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set MONITOREDSEVERGroup--+-----+----->><
                               '-group_name-'
```

Parameters

group_name

Specifies the IBM Spectrum Protect™ server group name that contains all monitored servers. You can remove a monitored server group name by issuing the command without specifying a value, or by specifying an empty value (""). Any existing monitoring for alerts and status from remote servers is ended.

Set the name of a monitored server group

Set the name of a monitored server group SUBS, by issuing the following command:

```
set monitoredservergroup subs
```

Remove the name of a monitored server group

Remove the monitored server group, by issuing the following command:

```
set monitoredservergroup
```

Related commands

Table 1. Commands related to SET MONITOREDSEVERGROUP

Command	Description
DEFINE SERVERGROUP (Define a server group)	Defines a new server group.
DEFINE GRPMEMBER (Add a server to a server group)	Defines a server as a member of a server group.
DELETE GRPMEMBER (Delete a server from a server group)	Deletes a server from a server group.
QUERY SERVERGROUP (Query a server group)	Displays information about server groups.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET MONITORINGADMIN (Set the name of the monitoring administrator)	Set the name of the monitoring administrator.

SET MONITORINGADMIN (Set the name of the monitoring administrator)

Use this command to set the name of the monitoring administrator that is used to connect to the servers in the monitored server group.

To display the name of the monitored server group, issue the QUERY MONITORSETTINGS command.

The administrator name that you specify must match the name of an existing administrator, otherwise the command fails.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set MONITORINGADMIN--+-+-----+----->>  
                        '-admin_name-'
```

Parameters

admin_name

Specifies administrator names. You can remove names by issuing the command without specifying a value, or by specifying an empty value ("").

Set the monitoring administrator name

Set the name of the monitoring administrator to MONADMIN, by issuing the following command:

```
set monitoringadmin monadmin
```

Remove the monitoring administrator name

Remove the monitoring administrator, by issuing the following command:

```
set monitoringadmin ""
```

Related commands

Table 1. Commands related to SET MONITORINGADMIN

Command	Description
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET MONITOREDSEVERGROUP (Set the group of monitored servers)	Set the group of monitored servers.

SET NODEATRISKINTERVAL (Specifies at-risk mode for an individual node)

Use this command to adjust the at-risk evaluation mode for an individual node.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax

```
>>---Set NODEATRISKINTERVAL--node_name----->  
  
>---TYPE---+---DEFAULT-----+-----><  
          +-BYPASSED-----+  
          '-CUSTOM--Interval---value-'
```

Parameters

node_name (Required)

Specifies the name of the client node that you want to update.

TYPE (Required)

Specifies the at-risk evaluation type. Specify one of the following values:

DEFAULT

Specifies that the node is evaluated with the same interval that was specified for the nodes classification by the SET STATUSATRISKINTERVAL command. The value is either system or applications, or VM, and is determined by the status monitor.

For example, you can specify `TYPE = DEFAULT`, which allows the status monitor to go ahead and classify the node automatically. Then the interval that is used, is the interval that was defined for that classification by the SET STATUSATRISKINTERVAL command.

BYPASSED

Specifies that the node is not evaluated for at-risk status by the status monitor. The at risk status is also reported as bypassed to the Operations Center.

CUSTOM

Specifies that the node is evaluated with the specified interval, rather than the interval that was specified by the SET STATUSATRISKINTERVAL command.

Interval

Specifies the amount of time, in hours, between client backup activity before the status monitor considers the client to be at risk. You can specify an integer in the range 6 - 8808. You must specify this parameter when `TYPE = CUSTOM`. You do not specify this parameter when `TYPE = BYPASSED` or `TYPE = DEFAULT`. The interval value for all client types is set to 24 at server installation.

Set node name to use a custom 90 day at-risk interval

Set the at-risk interval for a node named *fred* to 90 days.

```
set nodeatriskinterval fred type=custom interval=2160
```

Bypass the at-risk interval evaluation

Bypass the at-risk interval checking for a node named *bob*.

```
set nodeatriskinterval bob type=bypassed
```

Related commands

Table 1. Commands related to set nodeatriskinterval

Command	Description
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET VMATRISKINTERVAL (Specifies the at-risk mode for an individual VM filespace)	Sets the at-risk mode for a VM filespace
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
QUERY NODE (Query nodes)	Displays partial or complete information about one or more clients.

Command	Description
QUERY FILESPACE (Query one or more file spaces)	Displays information about data in file spaces that belong to a client.

SET PASSEXP (Set password expiration date)

Use this command to set the expiration period for administrator and client node passwords. You can either set a common password expiration period for all administrators and client node passwords or selectively set password expiration periods.

Restriction: The SET PASSEXP command does not apply to passwords that authenticate with an LDAP directory server.

You can override the SET PASSEXP setting for one or more nodes by using the REGISTER NODE or UPDATE NODE command with the PASSEXP parameter.

The NODE or ADMIN parameters must be specified to change the password expiration period for client nodes or administrators with selectively set password expiration periods. If you do not specify the NODE or ADMIN parameters, *all* client node and administrator passwords will use the new password expiration period. If you selectively set a password expiration period for a client node or administrator that does not already have a set password expiration period, it is not modified if you later set a password expiration for all users.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set PASSExp--days--+-+-----+----->
|           .-,----- . |
|           v              | |
|'-Node-----node_name--+'
>-----><
|           .-,----- . |
|           v              | |
|'-Admin-----admin_name--+'
>-----><
```

Parameters

days (Required)

Specifies the number of days that a password remains valid.

You can specify from 1 to 9999 if you do not specify the NODE or the ADMIN parameter. If you specify the NODE or the ADMIN parameter, you can specify from 0 to 9999. A value of 0 means that the password never expires. If a password expires, the server prompts for a new password when the administrator or client node contacts the server.

Node

Specifies the name of the node for which you are setting the password expiration period. To specify a list of nodes, separate the names with commas and no intervening spaces. This parameter is optional.

Admin

Specifies the name of the administrator whose password expiration period you would like to set. To specify a list of administrators, separate the names with commas and no intervening spaces. This parameter is optional.

Example: Set the administrator and client node password expiration

Set the administrator and client node password expiration period to 45 days.

```
set passexp 45
```

Example: Set an administrator's password expiration

Set the administrator LARRY's password expiration period to 120 days.

```
set passexp 120 admin=larry
```

Related commands

Table 1. Commands related to SET PASSEXP

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes that are associated with a client node.

SET PRODUCTOFFERING (Set the product offering that is licensed to your enterprise)

Use the SET PRODUCTOFFERING command to define the IBM Spectrum Protect™ product offering that is licensed to your enterprise.

The definition is used to determine whether automatic storage capacity measurement calculations are required and made available for use by the IBM® License Metric Tool (ILMT). Run this command only if you are using ILMT to determine license consumption.

For product offerings where automatic storage capacity measurement calculations are made available for use by ILMT, the parameter also defines which capacity measurement approach is used for those calculations.

The capacity measurement approach is defined by the licensing terms of your specific product offering. To determine the currently calculated storage capacity for your product offering, see Verifying license compliance.

The same storage capacity information is made available to ILMT on a weekly interval. After an applicable product offering is defined by using this command, IBM Spectrum Protect makes the current capacity calculation for that offering available to the ILMT. After the initial capacity calculation is made available to ILMT, IBM Spectrum Protect updates the value weekly.

Privilege class

To run this command, you must have system privilege.

Syntax

```
>>-SET PRODUCTOFFERING--product_offering-----<<
```

Parameters

product_offering (Required)

Specifies a product offering. The maximum length of the text string is 255 characters. The following options are available:

ENTry

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Entry. This product offering uses a Per Managed Server licensing metric. Capacity measurements for this product offering are not applicable.

DATARet

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect for Data Retention. Capacity measurements for this product offering are not calculated automatically or made available for use by ILMT.

BASIC

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect. This product offering uses a processor value unit (PVU) licensing metric. Capacity measurements for this product offering are not applicable.

EE

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Extended Edition. This product offering uses a PVU licensing metric. Capacity measurements for this product offering are not applicable.

SUIte

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEcloud

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite - IBM Cloud Object Storage Option. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEEntry

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite Entry. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEArchive

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite - Archive. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEProtectier

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite - ProtecTier. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEFrontend

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite - FrontEnd. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

SUITEENTRYFrontend

Specifies that the product offering licensed in your enterprise is IBM Spectrum Protect Suite Entry - FrontEnd. Capacity measurements for this product offering are calculated automatically and made available for use by ILMT.

CLEAR

No product offering is specified.

Example: Set the product offering to IBM Spectrum Protect (BASIC)

```
set productoffering BASIC
```

Related commands

Table 1. Commands related to SET PRODUCTOFFERING

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET QUERYSCHEDPERIOD (Set query period for polling client nodes)

Use this command to regulate how often client nodes contact the server to obtain scheduled work when it is running in the client-polling scheduling mode.

Each client can set its own retry period at the time its scheduler is started. You can use this command to override the value specified by all clients that can connect with the server.

If client nodes poll more frequently for schedules, the nodes receive changes to schedules more quickly. However, increased polling by the client nodes also increases network traffic.

You can issue the QUERY STATUS command to display the value for the period between schedule queries. At installation, IBM Spectrum Protect™ is configured so that each client node determines its own value for this setting.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set QUERYSHedperiod--++-----+----->><
                          '-hours-'
```

Parameters

hours

Specifies the maximum number of hours the scheduler on a client node waits between attempts to contact the server to obtain a schedule. This parameter is optional. You can specify an integer from 1 to 9999. If you do not specify a value for this parameter, each client determines its own value for this parameter.

Example: Set the polling period for all client nodes

Have all clients using the polling scheduling mode contact the server every 24 hours.

```
set queryschedperiod 24
```

Related commands

Table 1. Commands related to SET QUERYSCHEDPERIOD

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET SCHEDMODES	Specifies the central scheduling mode for the server.

SET RANDOMIZE (Set randomization of scheduled start times)

Use this command to set randomized start times within the startup window of each schedule for clients by using the client-polling scheduling mode. A startup window is the start time and duration during which a schedule must be initiated. A client-polling scheduling mode is a client/server communication technique where the client queries the server for work.

Each schedule has a window during which it can be run. To balance network and server load, the start times for clients can be scattered across that window. Use this command to specify the fraction of the window over which start times for clients are distributed.

The randomization occurs at the beginning of the window to allow time for retries, if necessary. When the scheduling mode is not set to polling, randomization does not occur if the client's first contact with the server is after the start time for the event.

You can issue the QUERY STATUS command to display the value for the schedule randomization percentage. At installation, the value is 25 percent.

Set the randomization percentage to a value greater than 0 to prevent communication errors. Communication errors can result from a large group of clients contacting the server simultaneously. If you do experience communication errors, you can increase the randomization percentage so that client contact is spread out. This decreases the chance for communication overload and failure.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set RANDomize--percent----->><
```


Parameters

percent (Required)

Specifies the percentage of the startup window over which the start times for individual clients are distributed. You can specify an integer from 0 to 50.

A value of 0 indicates that no randomization occurs and that all clients run schedules at the beginning of the startup windows.

A value of 50 indicates that clients are assigned start times that are randomly scattered across the first half of each startup window.

At installation, this value is 25, indicating that the first 25 percent of the window is used for randomization.

If you have specified DURUNITS=INDEFINITE in the DEFINE SCHEDULE command, the percentage is applied to a 24 hour period. For example, a value of 25 percent would result in a 6 hour window.

Example: Set randomization of scheduled start times

Set randomization to 50 percent.

```
set randomize 50
```

Related commands

Table 1. Commands related to SET RANDOMIZE

Command	Description
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET SCHEDMODES	Specifies the central scheduling mode for the server.

SET REPLRECOVERDAMAGED (Specify whether damaged files are recovered from a replication server)

Use this command to enable the system-wide recovery of damaged files from a target replication server. If this setting is turned on, the node replication process can be configured to detect damaged files on the source replication server and replace them with undamaged files from the target replication server.

The REPLRECOVERDAMAGED system parameter affects all file recovery processes across all replication processes for all nodes and file spaces. File recovery is possible only if the server software, Version 7.1.1 or later, is installed on the source and target replication servers, and if the node data was replicated before the file damage occurred.

To display the current setting, use the QUERY STATUS command.

When you install the server, the default setting is ON.

If you upgrade the server and no damaged files are detected, the default setting is ON.

If you upgrade the server and damaged files are detected, the parameter is set to OFF, and a message is issued to indicate that the recovery of damaged files is disabled. The OFF setting prevents the server from scanning database tables for damaged objects that can be recovered. Prevention of the scan is necessary in case many damaged files are detected. In that case, a scan can take a considerable amount of time, and should be scheduled when use of server resources is at a minimum. When you are ready to start the scan and recover damaged files, you must issue the SET REPLRECOVERDAMAGED command and specify the ON setting. After the server successfully completes the scan, the REPLRECOVERDAMAGED system parameter is set to ON.

The following table describes how the REPLRECOVERDAMAGED system parameter and other parameters affect the recovery of damaged, replicated files.

Table 1. Settings that affect the recovery of damaged files

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result
OFF	YES, NO, or not specified	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
OFF	ONLY	YES or NO	An error message is displayed because files cannot be recovered when the REPLRECOVERDAMAGED system parameter is set to OFF.
ON	YES	YES or NO	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	NO	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
ON	ONLY	YES or NO	Damaged files are recovered from the target replication server, but standard node replication does not occur.
ON	Not specified	YES	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	Not specified	NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

.-Set REPLRECOVERDamaged-----ON-----
>>+-----+----->>
'-Set REPLRECOVERDamaged-----+Off+-'
'-ON--'

```

Parameters

ON

Specifies that node replication is enabled to recover damaged files from a target replication server.

OFF

Specifies that node replication is not enabled to recover damaged files from a target replication server.

Example: Enable recovery of damaged files

To specify a system-wide setting that enables the server to recover damaged files from a target replication server, issue the following command:

set replrecoveredamaged on

Related commands

Table 2. Commands related to SET REPLRECOVERDAMAGED

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
UPDATE NODE	Changes the attributes that are associated with a client node.

SET REPLRETENTION (Set the retention period for replication records)

To maintain adequate information about replication processes, you can use this command to adjust the length of time that the source replication server retains replication records in its database. The SET REPLRETENTION command specifies the retention period for client-node replication records in the source replication-server database. You can use client node replication records to monitor running and completed processes.

A replication record is created when REPLICATE NODE command processing is started. By default, IBM Spectrum Protect™ retains client-node replication records for 30 calendar days. A calendar day consists of 24 hours, from midnight to midnight. For example, suppose that the retention period is two calendar days. If a replication process completes at 11:00 p.m. on day *n*, a record of that process is retained for 25 hours until midnight on day *n+1*. To display the retention period for replication records, issue the QUERY STATUS command on the source replication server.

Issue the SET REPLRETENTION command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set REPLREtention--.-30-----+----->>  
                        '-number_of_days-'
```

Parameters

number_of_days (Required)

The number of days that the source replication server retains replication records. You can specify an integer 0 - 9999. The default value is 30.

Example: Set a retention period for client-node replication records

You want to retain client-node replication records for 10 days.

```
set replretention 10
```

Related commands

Table 1. Commands related to SET REPLRETENTION

Command	Description
QUERY REPLICATION	Displays information about node replication processes.

Command	Description
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY REPLRULE	Displays information about node replication rules.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET REPLSERVER (Set the target replication server)

Use this command to set the name of a target replication server. You can also use this command to change or remove a target replication server.

Issue this command on the server that acts as a source for replicated data.

To display the name of a target replication server, issue the QUERY STATUS command on a source replication server.

Important:

- The server name that you specify with this command must match the name of an existing server definition. It must also be the name of the server to be used as the target replication server. If the server name specified by this command does not match the server name of an existing server definition, the command fails.
- Use care when you are changing or removing a target replication server. If you change a target replication server, replicated client-node data is sent to a different target replication server. If you remove a target replication server, client node data is not replicated.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set REPLSERVER--+-----+-----<<
                    '-target_server_name-'
```

Parameters

target_server_name

Specifies the name of the target replication server. The name that you specify must match the name of an existing server. The maximum length of a name is 64 characters.

To remove a target replication server, issue the command without specifying a value.

Note: If you do not want to continue replicating data, you can remove the node replication configuration after you remove the target replication server.

Example: Set a target replication server

The name of the server that you want to set as the target replication server is SERVER1.

```
set replserver server1
```

Related commands

Table 1. Commands related to SET REPLSERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.

Command	Description
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
UPDATE SERVER	Updates information about a server.
REMOVE REPLNODE	Removes a node from replication.
REMOVE REPLSERVER	Removes a server from replication.

SET RETRYPERIOD (Set time between retry attempts)

Use this command to set the number of minutes the scheduler on a client node waits between retry attempts after a failed attempt to contact the server or after a scheduled command fails to process.

Each client can set its own retry period at the time its scheduler program is started. You can use this command to override the values specified by all clients that can connect with the server.

This command is used in conjunction with the SET MAXCMDRETRIES command to regulate the period of time and the number of retry attempts to run a failed command.

You can issue the QUERY STATUS command to display the value for the period between retries. At installation, IBM Spectrum Protect™ allows each client to determine its own retry period.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set RETRYPeriod--+-----+----->>
                    '-minutes-'
```

Parameters

minutes

Specifies the number of minutes the scheduler on a client node waits between retry attempts after a failed attempt to contact the server or after a scheduled command fails to process. When setting the retry period, set a time period that permits more than one retry attempt within a typical startup window. You can specify an integer from 1 to 9999.

Example: Set a fifteen minute time period between retry attempts

Have the client scheduler retry failed attempts to contact the server or to process scheduled commands every fifteen minutes.

```
set retryperiod 15
```

Related commands

Table 1. Commands related to SET RETRYPERIOD

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.

SET SCHEDMODES (Select a central scheduling mode)

Use this command to determine how the clients communicate with the server to begin scheduled work. You must configure each client to select the scheduling mode in which it operates.

Use this command with the SET RETRYPERIOD command to regulate the time and the number of retry attempts to process a failed command.

You can issue the QUERY STATUS command to display the value for the scheduling mode supported. At installation, this value is ANY.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SCHEDMODEs---+ANY-----+-----><
      +-Polling--+
      '-PRompted-'
```

Parameters

ANY

Specifies that clients can run in either the client-polling or the server-prompted scheduling mode.

POLLing

Specifies that only the client-polling mode can be used. Client nodes poll the server at prescribed time intervals to obtain scheduled work.

PRompted

Specifies that only the server-prompted mode can be used. This mode is only available for clients that communicate with TCP/IP. Client nodes wait to be contacted by the server when scheduled work needs to be performed and a session is available.

Example: Restrict scheduled operations to clients using client-polling

Clients can run under both server-prompted and client-polling central scheduling. You want to temporarily restrict the scheduled operations to clients that use the client-polling mode. If you set the schedule mode to POLLING, the server discontinues prompting clients to run scheduled commands. This means that any client scheduler using the server-prompted mode waits until you set the schedule mode to ANY or PROMPTED.

```
set schedmodes polling
```

Related commands

Table 1. Command related to SET SCHEDMODES

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.

SET SCRATCHPADRETENTION (Set scratch pad retention time)

Use this command to set the amount of time for which scratch pad entries are retained.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SET SCRATCHPADRETENTION--days-----><
```

Parameters

days (Required)

Specifies the number of days that a scratchpad entry is retained after the last update to the scratchpad entry. You can enter an integer in the range 1 - 9999.

Example: Retain scratch pad entries for 367 days after they are updated

```
set scratchpadretention 367
```

Related commands

Table 1. Commands related to SET SCRATCHPADRETENTION

Command	Description
DEFINE SCRATCHPADENTRY	Creates a line of data in the scratch pad.
DELETE SCRATCHPADENTRY	Deletes a line of data from the scratch pad.
QUERY SCRATCHPADENTRY	Displays information that is contained in the scratch pad.
UPDATE SCRATCHPADENTRY	Updates data on a line in the scratch pad.

SET SERVERHLADDRESS (Set the high-level address of a server)

Use this command to set the high-level address (IP) of a server. IBM Spectrum Protect™ uses the address when you issue a DEFINE SERVER command with CROSSDEFINE=YES. You must use the SET SERVERHLADDRESS command for all automatic client deployments.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SERVERHladdress--ip_address-----><
```

Parameters

ip_address (Required)

Specifies a server high-level address as a numeric dotted decimal name or a host name. If a host name is specified, a server that can resolve the name to the dotted decimal form must be available.

Example: Set the high-level address of a server

Set the high-level address of HQ_SERVER to 9.230.99.66.

```
set serverhladdress 9.230.99.66
```

Related commands

Table 1. Command related to SET SERVERHLADDRESS

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERLLADDRESS	Specifies the low-level address of a server.

Command	Description
SET SERVERPASSWORD	Specifies the server password.

SET SERVERLLADDRESS (Set the low-level address of a server)

Use this command to set the low-level address of a server. IBM Spectrum Protect™ uses the address when you issue a DEFINE SERVER command with CROSSDEFINE=YES.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SERVERLLaddress--tcp_port-----><
```

Parameters

tcp_port (Required)

Specifies the low-level address of the server. Generally, this address is identical to the TCPPOINT option in the server option file of the server.

Example: Set the low-level address of a server

Set the low-level address of HQ_SERVER to 1500.

```
set serverlladdress 1500
```

Related commands

Table 1. Command related to SET SERVERLLADDRESS

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERPASSWORD	Specifies the server password.

SET SERVERNAME (Specify the server name)

Use this command to change the server name. When you install the IBM Spectrum Protect™ server, the name is set at installation to SERVER1.

Use the QUERY STATUS command to display the server name.

If you migrate from ADSM to IBM Spectrum Protect, the name is set to ADSM or the name last specified to ADSM with a SET SERVERNAME command.

Important:

- If this is a source server for a virtual volume operation, changing its name can impact its ability to access and manage the data it has stored on the corresponding target server.
- To prevent problems related to volume ownership, do not change the name of a server if it is a library client.

When changing the name of a server, be aware of the following additional restrictions:

- Windows clients use the server name to identify which passwords belong to which servers. Changing the server name after the clients are connected forces the clients to reenter the passwords.

- You must set unique names on servers that communicate with each other. On a network where clients connect to multiple servers, it is recommended that all of the servers have unique names.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SERVERname--server_name-----><
```

Parameters

server_name (Required)

Specifies the new server name. The name must be unique across a server network for enterprise event logging, enterprise configuration, command routing, or virtual volumes. The maximum length of the name is 64 characters.

Example: Name the server

Name the server WELLS_DESIGN_DEPT.

```
set servername wells_design_dept
```

Related commands

Table 1. Command related to SET SERVERNAME

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET SERVERPASSWORD (Set password for server)

Use this command to set the password for communication between servers to support enterprise administration and enterprise event logging and monitoring.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SERVERPAssword--password-----><
```

Parameters

password (Required)

Specifies a password for the server. Other servers must have the same password in their definitions of this server. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

Example: Set a server password

Set the password for HQ_SERVER to agave234.

```
set serverpassword agave234
```

Related commands

Table 1. Command related to SET SERVERPASSWORD

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.

SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)

Use this command to set the server replication rule for space-managed data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for space-managed data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.

You can specify a normal-priority replication rule or a high-priority replication rule. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain space-managed data and backup data. Replication of the space-managed data is a higher priority than the backup data. To prioritize the space-managed data, issue the SET SPREPLRULEDEFAULT command and specify the ALL_DATA_HIGH_PRIORITY replication rule. To prioritize the backup data, issue the SET BKREPLRULEDEFAULT command and specify the ALL_DATA replication rule for backup data. The ALL_DATA rule for backup data replicates backup data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SPREPLRuledefault--+-ALL_DATA-----+-----><
                               +-ALL_DATA_HIGH_PRIORITY--+
                               '-NONE-----'
```

Parameters

- ALL_DATA
Replicates space-managed data with a normal priority.
- ALL_DATA_HIGH_PRIORITY
Replicates space-managed data with a high priority.
- NONE
Space-managed data is not replicated.

Example: Set the server replication rule for space-managed data

Set up the default rule for space-managed data to replicate with a high priority.

```
set spreplruledefault all_data_high_priority
```

Related commands

Table 1. Commands related to SET BKREPLRULEDEFAULT

Command	Description
---------	-------------

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)

Use this command to adjust the backup activity interval that is used when the status monitor assesses whether clients are at risk.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>---Set STATUSATRISKINTERVAL--TYPE-------+-All-----+----->
                                     +-Applications--+
                                     +-VM-----+
                                     '-SSystems-----'
>----Interval--==value----->>
```

Parameters

TYPE (Required)

Specifies the type of client that should be evaluated. Specify one of the following values:

ALL

Specify this setting for all client types.

Applications

Specify this setting for only application client types.

VM

Specify this setting for virtual system clients types.

SSystems

Specify this setting for systems client types.

Interval (Required)

Specifies the amount of time, in hours, between client activity before the status monitor considers the client to be at risk. You can specify an integer in the range 6 - 8808. The interval value for all client types is set to 24 at server installation.

Set systems to use a two-week at-risk interval

Set the at-risk interval check for systems client types to 2 weeks.

```
set statusriskinterval type=systems interval=336
```

Related commands

Table 1. Commands related to

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

SET STATUSMONITOR (Specifies whether to enable status monitoring)

Use this command to enable and disable status monitoring. Turning status monitoring on for the first time also sets the default threshold values, and increases the event record retention to at least 14 days.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
.-Set STATUSMonitor-----Off-----.  
>>+-----+----->>  
'-Set STATUSMonitor-----+ON--+-'  
      '-Off-'
```

Parameters

ON

Specifies that the status monitoring is turned on. The first time that you set status monitoring to ON, it sets all the default threshold values that are specified in the DEFINE STATUSTHRESHOLD and UPDATE STATUSTHRESHOLD commands. It also sets the retention value for event records to at least 14 days. For example, when you turn status monitoring on, the default values for primary storage pool utilization is automatically set to display a warning when the threshold value reaches 80%, and an error when the threshold reaches 90% utilization.

OFF

Specifies that the status monitoring is turned off. Off is the default value.

Enable status monitoring

Set status monitoring to on to enable status monitoring.

```
set statusmonitor on
```

Related commands

Table 1. Commands related to SET STATUSMONITOR

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)

Use this command to specify the number of minutes between status monitoring server queries.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set STATUSREFreshinterval--minutes----->>
```

Parameters

minutes (Required)

Specifies the approximate number of minutes between status monitoring server queries. You can specify an integer in the range 1 - 2440. The default value is 5.

Restrictions:

- In a storage environment that is monitored by the Operations Center, set the same refresh interval on the hub and spoke servers. If you use different intervals, the Operations Center can show inaccurate information for spoke servers.
- Short status refresh intervals use more space in the server database and might require more processor and disk resources. For example, decreasing the interval by half doubles the required database and archive log space. Long intervals reduce the currency of Operations Center data but better suit a high-latency network configuration.
- A status refresh interval of less than 5 minutes can cause the following issues:

- Operations Center data that is supposed to be refreshed after the defined interval takes a longer time to be refreshed.
- Operations Center data that is supposed to be refreshed almost immediately when a related change occurs in the storage environment also takes a longer time to be refreshed.

Set the refresh interval for status monitoring

Specify that the server status is queried every 6 minutes, by issuing the following command:

```
set statusrefreshinterval 6
```

Related commands

Table 1. Commands related to SET STATUSREFRESHINTERVAL

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)

Use this command to enable the status monitor to consider clients as at risk when evaluating the status for each client.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set STATUSSKIPASFAILURE--+Yes-+----->
                               '-No--'
```

```
>--TYPE--++All-----><
          +-Applications-+
          +-VM-----+
          '-Systems-----'
```

Parameters

State (Required)

Specifies whether to enable the check for skipped files during the last backup. This check signifies that the client is at-risk if any files were skipped. Client data that is skipped or not backed up properly is considered at risk.

Yes

Specifies that the server evaluates whether a client is at risk.

No

Specifies that the server does not evaluate whether a client is at risk.

TYPE (Required)

Specifies the type of client that should be evaluated. Specify one of the following values:

ALL

Specify this setting for all client types.

APplications

Specify this setting for only application client types.

VM

Specify this setting for virtual system clients types.

SyStems

Specify this setting for systems client types.

Disable at-risk evaluation for virtual system client types

Disable the at-risk evaluation for virtual systems client types by issuing the following command:

```
set statusskipasfailure off type=vm
```

Related commands

Table 1. Commands related to SET STATUSSKIPASFAILURE

Command	Description
DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	Defines a status monitoring threshold.
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

SET SUBFILE (Set subfile backup for client nodes)

Use this command to set up the server to allow clients to back up subfiles. On the client's workstation, the SUBFILECACHEPATH and SUBFILECACHESIZE options must be specified in the client's options file (dsm.opt). If you are using a Windows client, you must also specify the SUBFILEBACKUP option.

With subfile backups, when a client's file has been previously backed up, any subsequent backups are typically made to the portion (a subfile) of the client's file that has changed, rather than the entire file.

Use the QUERY STATUS command to determine whether subfiles can be backed up to the server running this command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SUBFILE---Client+-----><
                '-No-----'
```

Parameters

Client

Specifies that the client node can determine whether to use subfile backup.

No

Specifies that the subfile backups are not to be used. At installation, this value is set to No.

Example: Set subfile backup for client nodes

Allow the client node to backup subfiles on the server.

```
set subfile client
```

Related commands

Table 1. Command related to SET SUBFILE

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)

Use this command to specify the number of days to keep information in the SQL activity summary table.

The SQL activity summary table contains statistics about each client session and server processes. For a description of the information in the SQL activity summary table, issue the following command:

```
select colname, remarks from columns where tablename='SUMMARY'
```

Issue the QUERY STATUS command to display the number of days the information is kept. At installation, IBM Spectrum Protect™ allows each server to determine its own number of days for keeping information in the SQL activity summary table.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-Set SUMmaryretention+-----><
                '-days-'
```

Parameters

days

Specifies the number of days to keep information in the activity summary table. You can specify a number from 0 to 9999. A value of 0 means that information in the activity summary table is not kept. A value of 1 specifies to keep the activity summary table for the current day.

Example: Specify the number of days to keep information in the SQL activity summary table

Set the server to retain the activity summary table information for 15 days.

```
set summaryretention 15
```

Related commands

Table 1. Commands related to SET SUMMARYRETENTION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.
QUERY ACTLOG	Displays messages from the server activity log.
SELECT	Allows customized queries of the IBM Spectrum Protect database.

SET TAPEALERTMSG (Set tape alert messages on or off)

Use this command to allow the IBM Spectrum Protect™ server to log notification of diagnostic information from library and drive devices. At installation, this value is set to OFF. When enabled, the server can retrieve diagnostic information from a tape or library device and display it using ANR messages. When disabled, the server will not query a device for this information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-Set TAPEAlertmsg--+-ON--+------>>  
          '-OFF-'
```

Parameters

ON

Specifies that diagnostic information will be reported to the server.

OFF

Specifies that diagnostic information will not be reported to the server.

Example: Set tape alert messages on

Allow the server to receive diagnostic information messages.

```
set tapealertmsg on
```

Related commands

Table 1. Command related to SET TAPEALERTMSG

Command	Description
---------	-------------

Command	Description
QUERY TAPEALERTMSG	Displays whether the server logs hardware diagnostic information.

SET TOCLOADRETENTION (Set load retention period for table of contents)

Use this command to specify the approximate number of minutes that unreferenced table of contents data will remain loaded in the server database.

During NDMP-controlled backup operations of NAS file systems, the server can optionally collect information about files and directories in the image and store this information in a table of contents within a storage pool. The web client can be used to examine files and directories in one or more file-system images by displaying entries from the table of contents data. The server loads the necessary table of contents data into a temporary database table.

Once the data have been loaded, the user can then select those files and directories to be restored. Because this database table is temporary, the data will only remain loaded for a specified time since the last reference to that data. At installation, the retention time is set to 120 minutes. Use the QUERY STATUS command to see the table of contents load retention time.

Privilege class

To issue this command you must have system privilege.

Syntax

```
>>-Set TOCLOADRetention--minutes-----<<
```

Parameters

minutes (Required)

Specifies the approximate number of minutes that an unreferenced table of contents data is retained in the database. You can specify an integer from 30 to 1000.

Example: Define the load retention period for the table of contents

Use the command, SET TOCLOADRETENTION, to specify that unreferenced table of contents data is to be retained in the database for 45 minutes.

```
set toclloadretention 45
```

Related commands

Table 1. Commands related to SET TOCLOADRETENTION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET VMATRISKINTERVAL (Specifies the at-risk mode for an individual VM filespace)

Use this command to adjust the at-risk evaluation mode for an individual VM filespace.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax

```
>>---Set VMATRISKINTERVAL--node_name--fsid----->>
>--TYPE---+--DEFAULT---+-------+----->>
      +-BYPASSED+  '-Interval---value-'
      '-CUSTOM---'
```

Parameters

node_name (Required)

Specifies the name of the client node, that owns the VM filespace, that you want to update.

fsid (Required)

Specifies the filespace ID of the client node that you want to update.

TYPE (Required)

Specifies which at-risk evaluation mode the status monitor should use when evaluating the at-risk classification for the specified nodes VM filespace. Specify one of the following values:

DEFAULT

Specifies that the VM filespace is evaluated with the same interval that was specified for the SET STATUSATRISKINTERVAL command.

BYPASSED

Specifies that the VM filespace is not evaluated for at-risk status by the status monitor. The at-risk status is also reported as bypassed to the Operations Center.

CUSTOM

Specifies that the VM filespace is evaluated with the specified interval, rather than the interval that was specified for the SET STATUSATRISKINTERVAL command.

Interval

Specifies the amount of time, in hours, between client backup activity before the status monitor considers the client to be at risk. You can specify an integer in the range 6 - 8808. You must specify this parameter when TYPE = CUSTOM. You do not specify this parameter when TYPE = BYPASSED or TYPE = DEFAULT. The interval value for all client types is set to 24 at server installation.

Set node name to use a custom 90 day at-risk interval

Set the at-risk interval for a node named *charlievm* (filespace ID 50) on datacenter node named *alice* to use a 90 day at-risk interval. You can issue the QUERY FILESPACE command to determine the filespace ID for the VM.

```
set vmatriskinterval alice 50 type=custom interval=2160
```

Bypass the at-risk interval evaluation

Exclude the VM called *davevm* (filespace ID 213) on datacenter node named *erin* from at-risk interval checking. You can issue the QUERY FILESPACE command to determine the filespace ID for the VM called *davevm*. Then set the at-risk interval check for the VM as bypassed.

```
set vmatriskinterval erin 213 type=bypassed
```

Related commands

Table 1. Commands related to set vmatriskinterval

Command	Description
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET NODEATRISKINTERVAL (Specifies at-risk mode for an individual node)	Sets the at-risk mode and interval for a node

Command	Description
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
QUERY NODE (Query nodes)	Displays partial or complete information about one or more clients.
QUERY FILESPACE (Query one or more file spaces)	Displays information about data in file spaces that belong to a client.

SETOPT (Set a server option for dynamic update)

You can use the SETOPT command to update most server options dynamically without stopping and restarting the server. For the DBDIAGLOGSIZE option, you must stop and start the server. A SETOPT command contained in a macro or a script cannot be rolled back.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-SETOPT--option_name--option_value-----<<
```

Parameters

option_name (Required)

Specifies a text string of information identifying the server option to be updated. The maximum length of the text string is 255 characters. The following options are available:

- ADMINCOMMTIMEOUT
- ADMINIDLETIMEOUT
- ALLOWREORGINDEX
- ALLOWREORGTABLE
- ARCHLOGCOMPRESS
- BACKUPINITIATIONROOT
- CHECKTAPEPOS
- CLIENTDEDUPTXNLIMIT
- CLIENTDEPLOYCATALOGURL
- CLIENTDEPLOYUSELOCALCATALOG
- COMMTIMEOUT
- Windows DATEFORMAT
- DBDIAGLOGSIZE
- DBDIAGPATHFSTHRESHOLD
- DEDUPTIER2FILESIZE
- DEDUPTIER3FILESIZE
- DEDUPREQUIRESBACKUP
- DNSLOOKUP
- EXPINTERVAL

- EXPQUIet
- FSUSEDThreshhold
- IDLETimeout
- LDAPCACHEDURATION
- MAXSessions
- MOVEBatchsize
- MOVESizethresh
- NDMPPREFDATAINTERFACE
- **Windows** NUMBERFORMAT
- NUMOPENVOLSallowed
- RECLAIMDELAY
- RECLAIMPERIOD
- REORGBEGINTime
- REORGDURation
- RESOURCETimeout
- RESTOREINTERVAL
- RETENTIONEXTENSION
- **AIX** | **Linux** | **Windows** SANDISCOVERY
- **AIX** | **Linux** | **Windows** SANREFRESHTIME
- SERVERDEDUPTXNlimit
- SHREDding
- **Windows** TCPPORT
- THROUGHPUTDatathreshold
- THROUGHPUTTimethreshold
- **Windows** TIMEFORMAT
- TXNGroupmax

option_value (Required)

Specifies the value for the server option.

Example: Set the maximum number of client sessions

Update the server option for the maximum number of client sessions to a value of 40.

```
setopt maxsessions 40
```

Related commands

Table 1. Commands related to SETOPT

Command	Description
QUERY OPTION	Displays information about server options.
QUERY SYSTEM	Displays details about the IBM Spectrum ProtectIBM Spectrum Protect™ server system.

SHRED DATA (Shred data)

Use this command to manually start the process of shredding deleted sensitive data. Manual shredding is possible only if automatic shredding is disabled.

You can control automatic shred processing with the SHREDDING server option.

This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

If data from a storage pool that enforces shredding is deleted while a manual shredding process is running, it will be added to the running process.

Privilege class

To issue this command you must have system privilege.

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background processes.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.

SUSPEND EXPORT (Suspend a currently running export operation)

Use this command to suspend a currently running server-to-server export operation which has a FILEDATA value that is not NONE. The export operation that you want to suspend must be past the initialization phase to be eligible for suspension. The state of the export operation is saved. The operation can be restarted by issuing the RESTART EXPORT command.

Privilege class

You must have system privilege to issue this command.

Syntax

```
>>-SUSPend EXPORT .*-----
+-----+-----><
'---export_identifier---
```

Parameters

EXPORTIdentifier

This optional parameter specifies the name of the export operation. You can find a name by issuing the QUERY EXPORT command to list all the currently running server-to-server export operations that can be suspended. You can also use the wildcard character to specify the name.

Example: Suspend a specific export operation

Suspend the running export operation EXPORTALLACCTNODES. No output is generated when you issue the SUSPEND EXPORT command. You must issue the QUERY EXPORT command to verify that the EXPORTALLACCTNODES operation is suspended.

```
suspend export exportallacctnodes
```

Example: Suspend all running export operations

Suspend all the export operations with a state of RUNNING.

```
suspend export *
```

Related commands

Table 1. Commands related to SUSPEND EXPORT

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
RESTART EXPORT	Restarts a suspended export operation.

UNLOCK commands

Use the UNLOCK commands to reestablish access after an object was locked.

- UNLOCK ADMIN (Unlock an administrator)
- UNLOCK NODE (Unlock a client node)
- UNLOCK PROFILE (Unlock a profile)

UNLOCK ADMIN (Unlock an administrator)

Use the UNLOCK ADMIN command to allow a locked administrator to access the server again. You can also unlock multiple administrators that authenticate with the same method.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UNLOCK Admin--+-*-----+--+-----+--><
                    '-admin_name-'   '-AUTHentication-----LOcal--'
                                      '-LDap--'
```

Parameters

admin_name (Required)

Specifies the name of the administrator to unlock. You can use wildcard characters to specify the administrator name. You do not have to enter an administrator name if you want to unlock all of the administrators according to their method of authentication. Use the wildcard with an authentication method to unlock multiple administrators. The parameter is required (no default wildcard).

AUTHentication

Specifies the method of password authentication that is needed for an administrator to log on.

LOcal

Specifies that you want to unlock administrator user IDs that authenticate passwords with the IBM Spectrum Protect™ server.

LDap

Specifies that you want to unlock administrator user IDs that authenticate passwords with an LDAP directory server.

Example: Unlock an administrator user ID

The administrator user ID JOE is locked out of IBM Spectrum Protect. Allow JOE to access the server. Issue the following command:

```
unlock admin joe
```

Example: Unlock all administrator user IDs that authenticate passwords with an LDAP directory server

The administrator user ID that use passwords that authenticate with an LDAP directory server must be unlocked so the IDs can communicate with the IBM Spectrum Protect server.

```
unlock admin * authentication=ldap
```

Related commands

Table 1. Commands related to UNLOCK ADMIN

Command	Description
---------	-------------

UNLOCK PROFILE (Unlock a profile)

Use this command on a configuration manager to unlock a configuration profile so it can be distributed to subscribing managed servers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UNLOCK PROFILE--profile_name-----<<
```

Parameters

profile_name (Required)

Specifies the profile to unlock. You can use wildcard characters to indicate multiple names.

Example: Unlock a profile

Unlock a profile named TOM.

```
unlock profile tom
```

Related commands

Table 1. Commands related to UNLOCK PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UPDATE PROFILE	Changes the description of a profile.

UPDATE commands

Use the UPDATE command to modify one or more attributes of an existing IBM Spectrum Protect™ object.

- UPDATE ADMIN (Update an administrator)
- UPDATE ALERTTRIGGER (Update a defined alert trigger)
- UPDATE ALERTSTATUS (Update the status of an alert)
- UPDATE BACKUPSET (Update a retention value assigned to a backup set)
- UPDATE CLIENTOPT (Update a client option sequence number)
- UPDATE CLOPTSET (Update a client option set description)
- UPDATE COLLOGROUP (Update a collocation group)
- UPDATE COPYGROUP (Update a copy group)
- UPDATE DATAMOVER (Update a data mover)
- UPDATE DEVCLASS (Update the attributes of a device class)
- UPDATE DOMAIN (Update a policy domain)

Specifies the category type for the alert, which is determined by the message types. The default value is SERVER.
 Note: Changing the category of an alert trigger does not change the category of existing alerts on the server. New alerts are categorized with the new category.
 Specify one of the following values:

Application

Alert is classified as application category. For example, you can specify this category for messages that are associated with application (TDP) clients.

Inventory

Alert is classified as inventory category. For example, you can specify this category for messages that are associated with the database, active log file, or archive log file.

Client

Alert is classified as client category. For example, you can specify this category for messages that are associated with general client activities.

Device

Alert is classified as device category. For example, you can specify this category for messages that are associated with device classes, libraries, drives, or paths.

Server

Alert is classified as general server category. For example, you can specify this category for messages that are associated with general server activities or events.

Storage

Alert is classified as storage category. For example, you can specify this category for messages that are associated with storage pools.

Systems

Alert is classified under system clients category. For example, you can specify this category for messages that are associated with system backup and archive or hierarchical storage management (HSM) backup-archive clients.

VMclient

Alert is classified under VMclient category. For example, you can specify this category for messages that are associated with virtual machine clients.

Admin

This optional parameter specifies the name of the administrator who receives email notification of this alert. The alert trigger is defined successfully even if no administrator names are specified.

ADDadmin

Specifies the administrator name that you want to add to the list of administrators that receive email alerts. Specify multiple administrator names, which are separated by commas, and no intervening spaces.

DELadmin

Specifies the administrator name that you want to delete from the list of administrators that receive email alerts. Specify multiple administrator names, which are separated by commas, and no intervening spaces.

Update alert trigger

Add the names of the administrators that want to be notified when ANR1073E, ANR1074E alerts occur, and also delete the name of an administrator that no longer wants to be notified, by issuing the following command:

```
update alertrigger ANR1073E,ANR1074E ADDadmin=djee,cdawson,mhaye deladmin=harryh
```

Related commands

Table 1. Commands related to UPDATE ALERTTRIGGER

Command	Description
DEFINE ALERTTRIGGER (Define an alert trigger)	Associates specified messages to an alert trigger.
DELETE ALERTTRIGGER (Remove a message from an alert trigger)	Removes a message number that can trigger an alert.
QUERY ALERTSTATUS (Query the status of an alert)	Displays information about alerts that have been issued on the server.
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	Displays message numbers that trigger an alert.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.

Command	Description
UPDATE ALERTSTATUS (Update the status of an alert)	Updates the status of a reported alert.

UPDATE ALERTSTATUS (Update the status of an alert)

Use this command to update the status of a reported alert.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-,------.
      v          |
>>-UPDate ALERTSStatus-----+--alert_id-+----->>
>--+-----+-----+-----+-----+----->
  '-SStatus-----+InActive-+-'  '-ASSigned-----text-'
      '-Closed---'
>--+-----+-----+-----+-----+-----><
  '-RESolvedby-----text-'  '-REMark-----text-'

```

Parameters

alert_id (Required)

Species the alert that you want to update. You can specify multiple message numbers by separating them with commas and no intervening spaces.

SStatus

Specifies the status type that you want to update. Alerts can be changed from active to inactive or closed, or from inactive to closed. Possible values are:

Inactive

Active alerts can be changed to inactive status.

Closed

Active and inactive alerts can be changed to closed status.

ASSigned

Specifies the administrator name that is assigned the alert that you want to query.

RESolvedby

Specifies the administrator name that resolved the alert that you want to query.

REMark

This parameter specifies comment text. The comment text cannot exceed 255 characters. If the description contains any blank spaces, enclose the entire text in quotation marks ("). Remove previously defined text by specifying a null string (") for this value.

Update the comment text in an alert

Issue the following command to update the comment text for alert ID number 25 and indicate that *DJADMIN* is working on the alert:

```
update alertstatus 25 assigned=DJADMIN
```

Update alert status

Issue the following command to change alert ID number 72 to the closed status, and add a remark about how the alert was resolved:

```
update alertstatus 72 status=closed remark="Increased the file system size for
the active log"
```

Related commands

Table 1. Commands related to UPDATE ALERTSTATUS

Command	Description
DEFINE ALERTTRIGGER (Define an alert trigger)	Associates specified messages to an alert trigger.
DELETE ALERTTRIGGER (Remove a message from an alert trigger)	Removes a message number that can trigger an alert.
QUERY ALERTSTATUS (Query the status of an alert)	Displays information about alerts that have been issued on the server.
QUERY ALERTTRIGGER (Query the list of defined alert triggers)	Displays message numbers that trigger an alert.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
UPDATE ALERTTRIGGER (Update a defined alert trigger)	Updates the attributes of one or more alert triggers.

UPDATE ADMIN (Update an administrator)

Use this command to change the password or contact information for an administrator. However, you cannot update the SERVER_CONSOLE administrator name.

AIX | **Linux** Passwords for administrators must be changed after a length of time that is determined by the SET PASSEXP command. The SET PASSEXP command does not affect passwords that authenticate with a Lightweight Directory Access Protocol (LDAP) server.

Restriction: You cannot update the authentication method for your own user ID. If necessary, another administrator must make that change. Also, when you update a password with the UPDATE ADMIN command, you cannot use a wildcard with the `admin_name` parameter.

Administrators with the same name as a node can be created during a REGISTER NODE command. To keep the node and administrator with the same name synchronized, the authentication method and the SSLREQUIRED setting for the node are updated to match the administrator. If the administrator authentication method is changed from LOCAL to LDAP and a password is not provided, the node is put in "LDAP pending" status. A password is then requested at the next logon. Passwords between same-named nodes and administrators are kept in sync through any authentication change.

You must use the RENAME ADMIN command to change the name of a registered administrator.

For users of Lightweight Directory Access Protocol (LDAP) servers:

- The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.
- If an administrative user ID matches a node name, do not update the authentication method to LDAP. If you do, you might see unexpected behavior because of automatic password changes that update the same password twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update operation might fail.

Privilege class

To issue this command to change another administrator password or contact information, you must have system privilege. Any administrator can issue this command to update his or her own password or contact information.

Syntax

```
>>-UPDate Admin-----admin_name-----+-----+----->
                                     '-password-'
>--+-----+-----+-----+-----+----->
   '-PASSExp----days-'   '-CONTACT----text-'
>--+-----+-----+-----+-----+----->
```


- For administrative user IDs that authenticate with an LDAP server, password expiration is set by using LDAP server utilities. For this reason, do not specify FORCEPWRESET=YES if you plan to specify AUTHENTICATION=LDAP.
- If you plan to update an administrative user ID to authenticate with an LDAP server, and you specified FORCEPWRESET=YES, you must change the password before you can specify FORCEPWRESET=NO and AUTHENTICATION=LDAP.

EMAILAddress

This parameter is used for additional contact information. The information that is specified by this parameter is not acted upon by IBM Spectrum Protect.

AUTHentication

This parameter determines the password authentication method that the administrator ID uses; either LDAP or LOCAL.

Local

Specifies that the administrator uses the local IBM Spectrum Protect server database to store passwords for authentication.

LDap

Specifies that the administrator uses an LDAP directory server for password authentication.

SYNCLdapdelete

This parameter applies only if an administrator who authenticates to an LDAP server wants to revert to local authentication.

Yes

Specifies that the administrator is deleted from the LDAP server.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Specifies that the administrator is not deleted from the LDAP server. This is the default.

SSLrequired (deprecated)

Specifies whether the administrator user ID must use the Secure Sockets Layer (SSL) protocol to communicate between the IBM Spectrum Protect server and the backup-archive client. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Important: Beginning with IBM Spectrum Protect Version 8.1.2 software and Tivoli® Storage Manager Version 7.1.8 software, this parameter is deprecated. Validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The SSLREQUIRED parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

SESSIONSECurity

Specifies whether the administrator must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRict

Specifies that the strictest security settings are enforced for the administrator. The STRICT value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the server and the administrator. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option.

To use the STRICT value, the following requirements must be met to ensure that the administrator can authenticate with the server:

- Both the administrator and server must be using IBM Spectrum Protect software that supports the SESSIONSECURITY parameter.
- The administrator must be configured to use the TLS 1.2 protocol for SSL sessions between the server and the administrator.

Administrators set to STRICT that do not meet these requirements are unable to authenticate with the server.

TRANSitional

Specifies that the existing security settings are enforced for the administrator. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the STRICT value.

If SESSIONSECURITY=TRANSITIONAL and the administrator has never met the requirements for the STRICT value, the administrator will continue to authenticate by using the TRANSITIONAL value. However, after an administrator meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. Then, the administrator can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT. In addition, after an administrator successfully authenticates by using a more secure communication protocol, the administrator can no longer authenticate by using a less secure protocol. For example, if an administrator that is not using SSL is updated and successfully authenticates by using TLS 1.2, the administrator can no longer authenticate by using no SSL protocol or TLS 1.1. This restriction also applies when you use functions such as command routing or server-to-server export, when the administrator authenticates to the IBM Spectrum Protect server as an administrator from another server.

ALert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No

Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the QUERY MONITORSETTINGS command.

Example: Update a password and password expiration period

Update the administrator LARRY to have the password SECRETWORD and a password expiration period of 120 days. The administrator in this example is authenticated to the IBM Spectrum Protect server.

```
update admin larry secretword passexp=120
```

Example: Update all administrators to communicate with a server by using strict session security

Update all administrators to use the strictest security settings to authenticate with the server.

```
update admin * sessionsecurity=strict
```

Related commands

Table 1. Commands related to UPDATE ADMIN

Command	Description
QUERY ADMIN	Displays information about one or more IBM Spectrum Protect administrators.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RENAME ADMIN	Changes an IBM Spectrum Protect administrator's name.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UPDATE NODE	Changes the attributes that are associated with a client node.

Related tasks:

Naming Tivoli Storage Manager objects

Related reference:

[Ssl client option](#)

UPDATE BACKUPSET (Update a retention value assigned to a backup set)

Use this command to update the retention value associated with a client's backup set.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax

```
.-,------.
v          |
>>-UPDate BACKUPSET-----+--node_name-----+----->
          '-node_group_name-'

.-,------.
v          |
>---backup_set_name+---REtention---+--days---+----->
          '-NOLimit-'

>--+-----+-----+-----+----->
  '-BEGINdate---date-' '-BEGINTime---time-'

>--+-----+-----+-----+----->
  '-ENDDate---date-' '-ENDTime---time-'

>--+-----+-----+-----+----->
  '-WHEREREtention---+--days---+-'
          '-NOLimit-'

.-WHEREDATAType---ALL------.
>--+-----+-----+-----+----->
|          .-,------.|
|          v          | |
|'-WHEREDATAType---+--FILE---+----->
|          '-IMAGE-'

>--+-----+-----+-----+----->
  '-WHEREDEsCRIPTION---description-'

.-VERSion---Any------.
>--+-----+-----+-----+----->>
  '-Preview---+--No---+' '-VERSion---+--Any---+'
          '-Yes-'          '-Latest-'
```

Parameters

node_name or **node_group_name** (Required)

Specifies the names of the client nodes or node groups whose data is contained in the specified backup set to be updated. To specify multiple node and node group names, separate the names with commas and no intervening spaces. The node names that you specify can contain wildcard characters, but node group names cannot contain wildcard characters.

backup_set_name (Required)

Specifies the name of the backup set to update. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

REtention (Required)

Specifies the updated number of days to retain the backup set on the server. You can specify an integer from 0 to 30000. The values are:

days

Specifies the updated number of days to retain the backup set.

NOLimit

Specifies that the backup set is retained on the server indefinitely. If you specify NOLIMIT, the server retains the volumes containing the backup set forever, unless a user or administrator deletes the volumes from server storage. Attention: Updating the retention period of a backup set may cause it to expire at a different time from other backup sets that might be stored on the same output media. In either case, the media will not be made available for other

uses until all of its backup sets have expired.

BEGINDate

Specifies the beginning date in which the backup set to update was created. This parameter is optional. The default is the current date. You can use this parameter with the BEGINTIME parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days <i>or</i> +days	The current date plus days specified.	TODAY +3 <i>or</i> +3.
TODAY-days <i>or</i> -days	The current date minus days specified.	TODAY-3 <i>or</i> -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time in which the backup set to update was created. This parameter is optional. The default is the current time. You can use this parameter with the BEGINDATE parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM <i>or</i> +HH:MM	The current time plus hours and minutes on the specified end date	NOW+02:00 <i>or</i> +02:00.
NOW-HH:MM <i>or</i> -HH:MM	The current time minus hours and minutes on the specified end date	NOW-02:00 <i>or</i> -02:00.

ENDDate

Specifies the ending date in which the backup set to update was created. This parameter is optional. You can use this parameter with the ENDTIME parameter to specify a range for the date and time. If you specify an end date without an ending time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+days <i>or</i> +days	The current date plus days specified.	TODAY +3 <i>or</i> +3.
TODAY-days <i>or</i> -days	The current date minus days specified.	TODAY -3 <i>or</i> -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

Value	Description	Example
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time in which the backup set to update was created. This parameter is optional. You can use this parameter with the ENDDATE parameter to specify a range for the date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00.

WHERERetention

Specifies the retention value, specified in days, that is associated with the backup set to update. The values are:

days

Specifies that the backup set that is retained this number of days is updated.

NOLimit

Specifies that the backup set retained indefinitely is updated.

WEREDescription

Specifies the description that is associated with the backup set to update. This parameter is optional. You can specify wildcard characters for the description. Enclose the description in quotation marks if it contains any blank characters.

WHEREDataType

Specifies the backup sets containing the specified types of data are to be updated. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be updated. To specify multiple data types, separate each data type with a comma and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be updated. This is the default.

FILE

Specifies that a file level backup set is to be updated. File level backup sets contain files and directories backup up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be updated. Image backup sets contain images created by the backup-archive client BACKUP IMAGE command.

Preview

Specifies whether to preview the list of backup sets to update, without actually updating the backup sets. This parameter is optional. The default is No. The values are:

No

Specifies that the backup sets are updated.

Yes

Specifies that the server displays the backup sets to update, without actually updating the backup sets.

VERsion

Specifies the version of the backup set to update. Backup sets with the same prefix name are considered to be different versions of the same backup set. This parameter is optional. The default is to update any version that matches the criteria specified on the command. The values are:

Any

Specifies that any version that matches the criteria specified on the command should be updated.

Latest

Specifies that only the most recent version of the backup set should be updated. If other criteria specified on the command (for example, ENDDATE or WHERERETENTION) exclude the most recent version of the backup set, then no backup set will be updated.

Example: Update a retention period

Update the retention period where the description is Healthy Computers. The retention period is assigned to backup set PERS_DATA.3099 that contains data from client node JANE. Change the retention period to 70 days.

```
update backupset jane pers_data.3099
retention=70 wheredescription="healthy computers"
```

Related commands

Table 1. Commands related to UPDATE BACKUPSET

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Updates a retention value associated with a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE NODEGROUP	Updates the description of a node group.

UPDATE CLIENTOPT (Update a client option sequence number)

Use this command to update the sequence number of a client option in a client option set.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax

```
>>-UPDate CLIENTOpt--option_set_name--option_name----->
>--current_sequence_number--new_sequence_number-----><
```

Parameters

- option_set_name (Required)
Specifies the name of the option set.
- option_name (Required)
Specifies a valid client option.
- current_sequence_number (Required)
Specifies the current sequence number of the option.
- new_sequence_number (Required)
Specifies the new sequence number of the option.

Example: Update a client option sequence number

To update the current client option sequence number issue the following command:

```
update clientopt eng dateformat 0 9
```

Related commands

Table 1. Commands related to UPDATE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.

UPDATE CLOPTSET (Update a client option set description)

Use this command to update the description for a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
>>-UPDate CLOptset--option_set_name----->
>--DEScRiption--===description-----<<
```

Parameters

- option_set_name (Required)
Specifies the name of the option set.
- DEScRiption (Required)
Specifies a description of the client option set. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

Example: Update a client option set description

Update the description for a client option set named ENG.

```
update cloptset eng description="unix"
```

Related commands

Table 1. Commands related to UPDATE CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.

UPDATE COLLOGROUP (Update a collocation group)

Use this command to modify the description of a collocation group.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
>>-UPDate COLLOGGroup--group_name----->
>--DESCRiption===--description-----><
```

Parameters

group_name

Specifies the name of the collocation group whose description you want to update.

DESCRiption (Required)

Specifies a description of the collocation group. This parameter is required. The maximum length of the description is 255 characters. If the description contains any blanks, enclose the entire description in quotation marks.

Example: Update a collocation group

Update the collocation group, GROUP1, with a new description.

```
update collogroup group1 "Human Resources"
```

Related commands

Table 1. Commands related to UPDATE COLLOGROUP

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE STGPOOL	Changes the attributes of a storage pool.

UPDATE COPYGROUP (Update a copy group)

Use this command to update a backup or archive copy group. To allow clients to use the updated copy group, you must activate the policy set that contains the copy group.

Tip: The UPDATE COPYGROUP command fails if you specify a copy storage pool as a destination.

The UPDATE COPYGROUP command takes two forms, depending upon whether the update is for a backup copy group or for an archive copy group. The syntax and parameters for each form are defined separately.

Table 1. Commands related to UPDATE COPYGROUP

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.

- UPDATE COPYGROUP (Update a backup copy group)
Use this command to update a defined backup copy group.
- UPDATE COPYGROUP (Update a defined archive copy group)
Use this command to update a defined archive copy group.

UPDATE COPYGROUP (Update a backup copy group)

Use this command to update a defined backup copy group.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax


```

>>-UPDate COpYgroup--domain_name--policy_set_name--class_name--->
>--+-----+-----+-----+-----+-----+-----+----->
  '-STANDARD-' '-Type---Backup-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-DESTination---pool_name-' '-FREQuency---days-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-VERExists---number--+'
                        '-NOLimit-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-VERDeleted---number--+'
                        '-NOLimit-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-RETEExtra---days---+' '-RETOOnly---days---+'
                        '-NOLimit-'                        '-NOLimit-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-MODE---MODified--+'
                        '-ABSolute-'
>--+-----+-----+-----+-----+-----+-----+----->
  '-SERialization---SHRStatic---+'
                        +-Static-----+
                        +-SHRDYnamic--+
                        '-DYnamic----+'
>--+-----+-----+-----+-----+-----+-----+-----><
  '-TOCDestination---pool_name---'

```

Parameters

domain_name (Required)

Specifies the policy domain to which the copy group belongs.

policy_set_name (Required)

Specifies the policy set to which the copy group belongs. You cannot update a copy group in the ACTIVE policy set.

class_name (Required)

Specifies the management class to which the copy group belongs.

STANDARD

Specifies the copy group, which must be STANDARD. This parameter is optional.

Type=Backup

Specifies that you want to update a backup copy group. This parameter is optional.

DESTination

Specifies the primary storage pool where the server initially stores backup data. This parameter is optional. You cannot specify a copy storage pool as the destination.

FREQuency

Specifies how frequently the server can back up a file. This parameter is optional. The server backs up a file only when the specified number of days has elapsed since the last backup. The FREQUENCY value is used only during a full incremental backup operation. This value is ignored during selective backup or partial incremental backup. You can specify an integer from 0 to 9999. The value 0 means that the server can back up a file regardless of when the file was last backed up.

VERExists

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional.

If an incremental backup causes the limit to be exceeded, the server expires the oldest backup version that exists in server storage. Possible values are:

number

Specifies the number of backup versions to retain for files that are currently on the client file system. You can specify an integer from 1 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 2. Preferred values are 3, 4, or more.

NOLimit

Specifies that you want the server to retain all backup versions.

The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using the server. This parameter is optional.

If a user deletes a file from the client file system, the next incremental backup causes the server to change the active backup version of the file to inactive and expire the oldest versions in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter. Possible values are:

number

Specifies the number of backup versions to retain for files that are deleted from the client file system after being backed up. You can specify a value from 0 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEExtra

Specifies the number of days that the server retains a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes inactive versions based on retention time even if the number of inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. Possible values are:

days

Specifies the number of days to retain inactive backup versions. You can specify an integer from 0 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 14 days. The preferred value is 30 or more days.

NOLimit

Specifies that you want to retain inactive backup versions indefinitely.

If you specify NOLIMIT, the server deletes extra backup versions based on the VEREXISTS parameter (when the file still exists on the client file system) or the VERDELETED parameter (when the file no longer exists on the client file system).

RETOOnly

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. Possible values are:

days

Specifies the number of days to retain the last remaining inactive copy of a file. You can specify an integer from 0 to 9999.

Tip: To help ensure that files can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 30 days.

NOLimit

Specifies that you want to keep the last remaining inactive version of a file indefinitely.

If you specify NOLIMIT, the server retains the last remaining backup version forever, unless a user or administrator deletes the file from server storage.

MODE

Specifies whether the server backs up a file only if the file has changed since the last backup, or whenever a client requests a backup. This parameter is optional. Possible values are:

MODified

Specifies that the file is backed up only if it has changed since the last backup. A file is considered changed if any of the following is true:

- The date last modified is different
- The file size is different
- The file owner is different
- The file permissions are different

ABSolute

Specifies that the file is backed up regardless of whether it has been changed.

The MODE value is used only for full incremental backup. This value is ignored during partial incremental backup or selective backup.

SERialization

Specifies how the server processes files or directories when they are modified during backup processing. This parameter is optional. Possible values are:

SHRStatic

Specifies that the server backs up a file or directory only if it is not being modified during backup. The server attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file or directory is modified during each backup attempt, the server does not back it up.

Static

Specifies that the server backs up a file or directory only if it is not being modified during backup. The server attempts to perform the backup only once.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file or directory is being modified during a backup attempt, the server backs up the file or directory during the last attempt even though the file or directory is being modified. The server attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that the server backs up a file or directory on the first attempt, regardless of whether the file or directory is being modified during backup processing.

Important: Be careful about using the SHR DYNAMIC and DYNAMIC values. IBM Spectrum Protect™ uses these values to determine if it backs up a file or directory while modifications are occurring. As a result, the backup version might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file or directory because it contains some, but not all, modifications. If a file that contains a fuzzy backup is restored, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that IBM Spectrum Protect creates a backup version only if the file or directory is not being modified.

TOCDestination

Specifies the primary storage pool in which a table of contents (TOC) will initially be stored for any NDMP backup or backup set operation for which a TOC is generated. This parameter is optional. You cannot specify a copy storage pool as the destination. The storage pool specified for the destination must have NATIVE or NONBLOCK data format. To avoid mount delays, ensure that the storage pool has a device class of DISK or DEVTYPE=FILE. TOC generation is an option for NDMP backup operations, but is not supported for other image-backup operations.

To remove an existing TOC destination from the copy group, specify a null string ("") for this value.

If TOC creation is requested for a backup operation that uses NDMP and the image is bound to a management class whose backup copy group does not specify a TOC destination, the outcome will depend on the TOC parameter for the backup operation.

- If TOC=PREFERRED (the default), the backup proceeds without creation of a TOC.
- If TOC=YES, the entire backup fails because no TOC can be created.

Example: Update a backup copy group

Update the backup copy group (STANDARD) in the EMPLOYEE_RECORDS policy domain, VACATION policy set, ACTIVEFILES management class. Change the destination to DISKPOOL, with a minimum interval of seven days between backups, regardless of whether the files have been modified. Retain up to three backup versions while a file still exists on a client file system.

```
update copygroup employee_records vacation
activefiles type=backup destination=diskpool
frequency=7 verexists=3 mode=absolute
```

UPDATE COPYGROUP (Update a defined archive copy group)

Use this command to update a defined archive copy group.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax

```
>>-UPDate CCopygroup--domain_name--policy_set_name--class_name--->
>--+-----+---Type---+---Archive----->
  '-STANDARD-'
>--+-----+-----+-----+----->
  '-DESTination---+pool_name-' '-FREQuency---+Cmd-'
>--+-----+-----+-----+----->
  '-RETVer---+days---+' '-MODE---+ABSolute-'
      '-NOLimit-'
>--+-----+----->
  '-RETMin---+days---'
>--+-----+----->>
  '-SERialization---+SHRStatic---+'
                        +-Static-----+
                        +-SHRDYnamic--+
                        '-DYnamic----'
```

Parameters

domain_name (Required)

Specifies the policy domain to which the copy group belongs.

policy_set_name (Required)

Specifies the policy set to which the copy group belongs. You cannot update a copy group in the ACTIVE policy set.

class_name (Required)

Specifies the management class to which the copy group belongs.

STANDARD

Specifies the copy group, which must be STANDARD. This parameter is optional.

Type=Archive (Required)

Specifies that you want to update an archive copy group. This parameter is required.

DESTination

Specifies the primary storage pool where the server initially stores the archive copy. This parameter is optional. You cannot specify a copy storage pool as the destination.

FREQuency=Cmd

Specifies the copy frequency, which must be CMD. This parameter is optional.

RETVer

Specifies the number of days to keep an archive copy. This parameter is optional. Possible values are:

days

Specifies the number of days to keep an archive copy. You can specify an integer from 0 to 30000.

Tip: To help ensure that your data can be recovered after a malware incident, such as a ransomware attack, specify a value of at least 30 days.

NOLimit

Specifies that you want to keep an archive copy indefinitely.

If you specify NOLIMIT, the server retains archive copies forever, unless a user or administrator deletes the file from server storage.

The value of the RETVER parameter can affect the management class to which the server binds an archived directory. If the client does not use the ARCHMC option, the server binds directories that are archived to the default management class. If the default management class has no archive copy group, the server binds directories that are archived to the management class with the shortest retention period.

MODE=ABSolute

Specifies that a file is always archived when the client requests it. The MODE must be ABSOLUTE. This parameter is optional.

RETMIn

Specifies the minimum number of days to keep an archive copy after it has been archived. This parameter is optional. The default value is 365.

SERialization

Specifies how the server processes files that are modified during archive. This parameter is optional. Possible values are:

SHRStatic

Specifies that the server does not archive a file that is being modified. The server attempts to perform an archive as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file is modified during the archive attempt, the server does not archive the file.

Static

Specifies that the server does not archive a file that is being modified. If a file is modified during the archive attempt, the server does not archive the file.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file is being modified during an archive attempt, the server archives the file during its last attempt even though the file is being modified. The server attempts to archive the file as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that the server archives a file on the first attempt, regardless of whether the file is being modified during archive processing.

Important: Be careful about using the SHR DYNAMIC and DYNAMIC values. IBM Spectrum Protect™ uses them to determine if it archives a file while modifications are occurring. As a result, the archive copy might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file because it contains some, but not all, modifications. If a file that contains a fuzzy backup is retrieved, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that IBM Spectrum Protect creates an archive copy only if the file is not being modified.

Tip: Be cautious when selecting retention values for primary storage pools that are of type RECLAMATIONTYPE=SNAPLOCK. Volumes in these types of storage pools cannot be deleted until after their retention dates have passed.

Example: Update multiple elements of a copy group

Update the archive copy group (STANDARD) in the EMPLOYEE_RECORDS policy domain, VACATION policy set, ACTIVEFILES management class. Change the destination to TAPEPOOL. Keep archive copies for 190 days.

```
update copygroup employee_records vacation
activefiles standard type=archive
destination=tapepool retver=190
```

UPDATE DATAMOVER (Update a data mover)

Use this command to update the definition for a data mover or set a data mover off-line when the hardware is being maintained.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DATAMover--data_mover_name----->
>--+-----+-----+----->
  '-HLAddress---address-' '-LLAddress---tcp_port-'
>--+-----+-----+----->
  '-USERid---userid-' '-PASsword---password-'
>--+-----+-----+-----><
  '-ONLine---+Yes-+-'
```

Parameters

data_mover_name (Required)

Specifies the name of the data mover.

HLAddress

Specifies either the new numerical IP address or the new domain name, which is used to access the NAS file server. This parameter is optional.

LLAddress

Specifies the new TCP port number to access the NAS file server for Network Data Management Protocol (NDMP) sessions. This parameter is optional.

USERid

Specifies the user ID for a user that is authorized to initiate an NDMP session with the NAS file server. For example, enter the administrative ID for a NetApp file server. This parameter is optional.

PASsword

Specifies the new password for the user ID to log onto the NAS file server. This parameter is optional.

ONLine

Specifies whether the data mover is available for use. This parameter is optional.

Yes

Specifies that the data mover is available for use.

No

Specifies that the data mover is not available for use.

Attention: If a library is controlled using a path from a data mover to the library, and the data mover is offline, the server will not be able to access the library. If the server is halted and restarted while the data mover is offline, the library will not be initialized.

Example: Update a data mover IP address

Update the data mover for the node named NAS1. Change the numerical IP address from 9.67.97.103 to 9.67.97.109.

```
update datamover nas1 hladdress=9.67.97.109
```

Example: Update a data mover domain name

Update the data mover for the node named NAS1. Change the numerical IP address from 9.67.97.109 to the domain name of NETAPP2.TUCSON.IBM.COM.

```
update datamover nas1 hladdress=netapp2.tucson.ibm.com
```

Related commands

Table 1. Commands related to UPDATE DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DATAMOVER	Deletes a data mover.
QUERY DATAMOVER	Displays data mover definitions.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE NODE	Changes the attributes that are associated with a client node.

UPDATE DEVCLASS (Update the attributes of a device class)

Use this command to update a defined device class.

Note: The DISK device class is predefined by IBM Spectrum Protect™ and cannot be modified with the UPDATE DEVCLASS command.

AIX | Linux If you are updating a device class for devices that are to be accessed through a z/OS® media server, see UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server).

The syntax and parameter descriptions are provided according to the device type. The syntax and parameter information is presented in the following order.

- UPDATE DEVCLASS (Update a 3590 device class)
- UPDATE DEVCLASS (Update a 3592 device class)
- UPDATE DEVCLASS (Update a 4MM device class)
- UPDATE DEVCLASS (Update an 8MM device class)
- UPDATE DEVCLASS (Update a CENTERA device class)
- UPDATE DEVCLASS (Update a DLT device class)
- UPDATE DEVCLASS (Update an ECARTRIDGE device class)
- UPDATE DEVCLASS (Update a FILE device class)
- **AIX | Windows** UPDATE DEVCLASS (Update a GENERICTAPE device class)
- UPDATE DEVCLASS (Update an LTO device class)
- UPDATE DEVCLASS (Update a NAS device class)
- UPDATE DEVCLASS (Update a REMOVABLEFILE device class)
- UPDATE DEVCLASS (Update a SERVER device class)
- UPDATE DEVCLASS (Update a VOLSAFE device class)

Table 1. Commands related to UPDATE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up IBM Spectrum Protect device information to a file.
DEFINE DEVCLASS	Defines a device class.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE LIBRARY	Changes the attributes of a library.

UPDATE DEVCLASS (Update a 3590 device class)

Use the 3590 device class when you are using 3590 tape devices.

AIX | Linux If you are defining a device class for devices that are to be accessed through a z/OS® media server, see UPDATE DEVCLASS (Update a 3590 device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-FORMAT---+--DRIVE---+'
                                     +-3590B---+
                                     +-3590C---+
                                     +-3590E-B--+
                                     +-3590E-C--+
                                     +-3590H-B--+
                                     '-3590H-C-'
>--+-----+----->
```


Format	Estimated Capacity	Description
3590H-C	See note 60.0 GB (J cartridge- standard length) 120.0 GB (K cartridge- extended length)	Compressed format, similar to the 3590C format
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be greater than the listed value.		

Table 2. 3590 device recording format selections

Device	Format					
	3590B	3590C	3590E-B	3590E-C	3590H-B	3590H-C
3590	Read/Write	Read/Write	–	–	–	–
Ultra-SCSI	Read/Write	Read/Write	–	–	–	–
3590E	Read	Read	Read/Write	Read/Write	–	–
3590H	Read	Read	Read	Read	Read/Write	Read/Write

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes that are categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

UPDATE DEVCLASS (Update a 3592 device class)

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see UPDATE DEVCLASS (Update a 3592 device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-Update DEVclass--device_class_name----->
>--+-----+----->
  '-LIBRARY---library_name-'
>--+-----+----->
  '-LBProtect---+READWrite+-'
                    +-WRITEOnly+
                    '-No-----'
>--+-----+-----+----->
  '-SCALECAPacity---+100-+-'  '-FORMAT---+DRIVE---+-'
                    +-90--+          +-3592-----+
                    '-20--'          +-3592C----+
                                       +-3592-2---+
                                       +-3592-2C--+
                                       +-3592-3---+
                                       +-3592-3C--+
                                       +-3592-4---+
```

```

+-3592-4C--+
+-3592-5---+
+-3592-5C--+
+-3592-5A--+
'-3592-5AC-'

>-----+----->
'-ESTCAPacity----size-'

>-----+----->
'-PREFIX----+ADSM-----+-'
      '-tape_volume_prefix-'

>-----+-----+----->
'-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'

>-----+----->
'-MOUNTLimit----+DRIVES--+-'
      +-number-+
      '-0-----'

>-----+-----><
| (1) (2) |
|-----DRIVEEncryption---+ON-----+|
      +-ALLOW-----+
      +-EXTERNAL-+
      '-OFF-----'

```

Notes:

1. You cannot specify both WORM=Yes and DRIVEENCRYPTION=ON.
2. Drive encryption is supported only for 3592 Generation 2 or later drives.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of the defined library object that contains the tape drives that can be used by this device class.

This parameter is optional.

For information about defining a library object, see the DEFINE LIBRARY command.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When LBPROTECT is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect™ and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the BACKUP DB command.

When the LBPROTECT parameter is set to READWRITE, you do not have to specify the CRCDATA parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to

generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the BACKUP DB command.

No

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on IBM® 3592 Generation 3 drives and later with 3592 Generation 2 media and later.

See Technote 1634851, Additional information on the Tivoli® Storage Manager LBProtect option, for an explanation about when to use the LBProtect parameter.

SCALECapacity

Specifies the percentage of the media capacity that can be used to store data. This parameter is optional. Possible values are 20, 90, or 100.

Setting the scale capacity percentage to 100 provides maximum storage capacity. Setting it to 20 provides fastest access time.

Note: The scale capacity value takes effect when data is first written to a volume. Any updates to the device class for scale capacity do not affect volumes that already have data that is written to them until the volume is returned to scratch status.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats, estimated capacities, and recording format options for 3592 devices.

Tip: The format name is specified as, for example, 3592-X, 3592-XC, 3592-XA, or 3592-XAC, where X indicates the drive generation, C indicates a compressed format, and A indicates an archive drive.

Table 1. Recording formats and default estimated capacities for 3592

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
3592	300 GB	Uncompressed (basic) format
3592C	See note.	Compressed format
3592-2	500 GB	Uncompressed (basic) format JA tapes
	700 GB	Uncompressed (basic) format JB tapes
3592-2C	1.5 TB	Compressed format JA tapes
	2.1 TB	Compressed format JB tapes
3592-3	640 GB	Uncompressed (basic) format JA tapes
	1 TB	Uncompressed (basic) format JB tapes
3592-3C	1.9 TB	Compressed format JA tapes
	3 TB	Compressed format JB tapes

Format	Estimated capacity	Description
3592-4	400 GB	Uncompressed (basic) format JK tapes
	1.5 TB	Uncompressed (basic) format JB tapes
	3.1 TB	Uncompressed (basic) format JC tapes
3592-4C	1.2 TB	Compressed format JK tapes
	4.4 TB	Compressed format JB tapes
	9.4 TB	Compressed format JC tapes
3592-5 (For IBM TS1150 Model 3592 E08 drives with product ID 03592E08)	900 GB	Uncompressed (basic) format JK tapes
	7 TB	Uncompressed (basic) format JC/JY tapes
	2 TB	Uncompressed (basic) format JL tapes
	10 TB	Uncompressed (basic) format JD/JZ tapes
3592-5C (For IBM TS1150 Model 3592 E08 drives with product ID 03592E08)	Depends on the compressibility of the data	Compressed format JK tapes
		Compressed format JC/JY tapes
		Compressed format JL tapes
		Compressed format JD/JZ tapes
3592-5A (For IBM TS1155 Model 3592 55F drives with product ID 0359255F)	3 TB	Uncompressed (basic) format JL tapes
	15 TB	Uncompressed (basic) format JD/JZ tapes
3592-5AC (For IBM TS1155 Model 3592 55F drives with product ID 0359255F)	Depends on the compressibility of the data	Compressed format JL tapes
		Compressed format JD/JZ tapes
Note: If this format uses the compression feature for tape drives, depending on the effectiveness of compression, the actual capacity might be different from the estimated capacity.		

Important: For optimal performance, avoid mixing different generations of drives in a single SCSI library.

Special configurations are also required for mixing different generations of 3592 drives in 349x and ACSLS libraries.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional.

Updating this parameter affects empty volumes only. If a filling volume was previously encrypted or is unencrypted, and you update the DRIVEENCRYPTION parameter, the volume maintains its original encrypted or unencrypted status. The filling volume also maintains its original key-management status.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes—for example, back up sets, export volumes, and database backup volumes—will not be encrypted.) If you specify ON and you enable either the library or system method of encryption, drive encryption is not allowed and backup operations fail.

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if either the library or system method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive.

When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption.

By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable either the library or system method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

UPDATE DEVCLASS (Update a 4MM device class)

Use the 4MM device class when you are using 4 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->>
>--+-----+-----+-----+----->
  '-LIBRARY---library_name-'  '-FORMAT-----DRIVE-+-'
                                     +-DDS1--+
                                     +-DDS1C-+
                                     +-DDS2--+
                                     +-DDS2C-+
                                     +-DDS3--+
                                     +-DDS3C-+
                                     +-DDS4--+
                                     +-DDS4C-+
                                     +-DDS5--+
                                     +-DDS5C-+
                                     +-DDS6--+
                                     '-DDS6C-'

>--+-----+-----+-----+----->
  '-ESTCAPacity---size-'

>--+-----+-----+-----+----->
  '-PREFIX-----ADSM-----+-'
      '-tape_volume_prefix-'

>--+-----+-----+-----+----->
  '-MOUNTWait---minutes-'  '-MOUNTRetention---minutes-'

>--+-----+-----+-----+----->>
  '-MOUNTLimit-----DRIVES-+-'
      +-number-+
      '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined.

LIBRARY

Specifies the name of the defined library object that contains the 4 mm tape drives used by this device class. This parameter is optional. For information about defining a library object, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for 4 mm devices:

Table 1. Recording formats and default estimated capacities for 4 mm tapes

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
DDS1	1.3 GB (60 meter) 2.0 GB (90 meter)	Uncompressed format, applies only to 60-meter and 90-meter tapes
DDS1C	See note 1.3 GB (60 meter) 2.0 GB (90 meter)	Compressed format, applies only to 60-meter and 90-meter tapes
DDS2	4.0 GB	Uncompressed format, applies only to 120-meter tapes
DDS2C	See note 8.0 GB	Compressed format, applies only to 120-meter tapes
DDS3	12.0 GB	Uncompressed format, applies only to 125-meter tapes
DDS3C	See note 24.0 GB	Compressed format, applies only to 125-meter tapes
DDS4	20.0 GB	Uncompressed format, applies only to 150-meter tapes
DDS4C	See note 40.0 GB	Compressed format, applies only to 150-meter tapes
DDS5	36 GB	Uncompressed format, when using DAT 72 media
DDS5C	See note 72 GB	Compressed format, when using DAT 72 media
DDS6	80 GB	Uncompressed format, when using DAT 160 media
DDS6C	See note 160 GB	Compressed format, when using DAT 160 media

Format	Estimated Capacity	Description
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be greater than the listed value.		

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes that are categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about the default estimated capacity for 4 mm tapes, see Table 1.

PREFIX

Specifies the high-level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

UPDATE DEVCLASS (Update an 8MM device class)

Use the 8MM device class when you are using 8 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+-----+-----+----->
  '-LIBRARY---library_name-' '-FORMAT---+DRIVE-+'
                                     +-8200--+
                                     +-8200C+
                                     +-8500--+
                                     +-8500C+
                                     +-8900--+
                                     +-AIT---+
                                     +-AITC--+
                                     +-M2----+
                                     +-M2C---+
                                     +-SAIT--+
                                     +-SAITC+
                                     +-VXA2--+
                                     +-VXA2C+
                                     +-VXA3--+
                                     '-VXA3C-'

>--+-----+-----+-----+----->
  '-ESTCAPacity---size-'

>--+-----+-----+-----+----->
  '-PREFIX---+ADSM-----+-'
                '-tape_volume_prefix-'

>--+-----+-----+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'

>--+-----+-----+-----+-----><
  '-MOUNTLimit---+DRIVES-+'
                +-number+
                '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the 8 mm tape drives that can be used by this device class.

For more information about defining a library object, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for 8 mm devices:

Table 1. Recording format and default estimated capacity for 8 mm tape

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
8200	2.3 GB	Uncompressed (standard) format, using standard 112-meter tape cartridges
8200C	See note 3.5 GB 4.6 GB	Compressed format, using standard 112-meter tape cartridges
8500	See note	Drives (Read Write)
15m	600 MB	Eliaint 820 (RW)
15m	600 MB	Exabyte 8500/8500C (RW)
15m	600 MB	Exabyte 8505 (RW)
54m	2.35 GB	Eliaint 820 (RW)
54m	2.35 GB	Exabyte 8500/8500C (RW)
54m	2.35 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliaint 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliaint 820 (RW)

Format		Description
Medium Type	Estimated Capacity	
8500C	See note	Drives (Read Write)
15m	1.2 GB	Eliant 820 (RW)
15m	1.2 GB	Exabyte 8500/8500C (RW)
15m	1.2 GB	Exabyte 8505 (RW)
54m	4.7 GB	Eliant 820 (RW)
54m	4.7 GB	Exabyte 8500/8500C (RW)
54m	4.7 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliant 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliant 820 (RW)
8900	See note	Drive (Read Write)
15m	–	Mammoth 8900 (R)
54m	–	Mammoth 8900 (R)
112m	–	Mammoth 8900 (R)
160m XL	–	Mammoth 8900 (R)
22m	2.5 GB	Mammoth 8900 (RW)
125m	–	Mammoth 8900 (RW with upgrade)
170m	40 GB	Mammoth 8900 (RW)
AIT	See note	Drive
SDX1–25C	25 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	35 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	36 GB	AIT2 and AIT3 drives
SDX2–50C	50 GB	AIT2 and AIT3 drives
SDX3–100C	100 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	150 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	200 GB	AIT4 and AIT5 drives
SDX5-400C	400 GB	AIT5 drive
AITC	See note	Drive
SDX1–25C	50 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	91 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	72 GB	AIT2 and AIT3 drives
SDX2–50C	130 GB	AIT2 and AIT3 drives
SDX3–100C	260 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	390 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	520 GB	AIT4 and AIT5 drives
SDX5-400C	1040 GB	AIT5 drive
M2	See note	Drive (Read Write)
75m	20.0 GB	Mammoth II (RW)
150m	40.0 GB	Mammoth II (RW)
225m	60.0 GB	Mammoth II (RW)
M2C	See note	Drive (Read Write)
75m	50.0 GB	Mammoth II (RW)
150m	100.0 GB	Mammoth II (RW)
225m	150.0 GB	Mammoth II (RW)
SAIT	See note	Drive (Read Write)
	500 GB	Sony SAIT1–500(RW)
SAITC	See note	Drive (Read Write)
	1300 GB (1.3 TB)	Sony SAIT1–500(RW)

Format		Description
Medium Type	Estimated Capacity	
VXA2	See note	Drive (Read Write)
V6 (62m)	20 GB	VXA-2
V10 (124m)	40 GB	
V17 (170m)	60 GB	
VXA2C	See note	Drive (Read Write)
V6 (62m)	40 GB	VXA-2
V10 (124m)	80 GB	
V17 (170m)	120 GB	
VXA3	See note	Drive (Read Write)
X6 (62m)	40 GB	VXA-3
X10 (124m)	86 GB	
X23 (230m)	160 GB	
VXA3C	See note	Drive (Read Write)
X6 (62m)	80 GB	VXA-3
X10 (124m)	172 GB	
X23 (230m)	320 GB	
<p>Note: The actual capacities might vary depending on which cartridges and drives are used.</p> <ul style="list-style-type: none"> • For the AITC and SAITC formats, the normal compression ratio is 2.6:1. • For the M2C format, the normal compression ratio is 2.5:1. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about the default estimated capacity for 8 mm tapes, see Table 1.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

Example: Update the mount limit and capacity of an 8 mm device class

Update a device class named 8MMTAPE. Change the mount limit to 3 and the estimated capacity to 10 GB.

```
update devclass 8mmtape mountlimit=3 estcapacity=10G
```

Example: Update the mount retention period of an 8 mm device class

Update an 8 mm device class that is named 8MMTAPE to a 15-minute mount retention.

```
update devclass 8mmtape mountretention=15
```

UPDATE DEVCLASS (Update a CENTERA device class)

Use the CENTERA device class when you are using EMC Centera storage devices. The CENTERA device type uses files as volumes to store data sequentially. It is similar to the FILE device class.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
                                     (1)
>--HLAddress---ip_address?PEA_file----->
>--+-----+----->
  '-MINCAPacity-----size---'
>--+-----+----->>
  '-MOUNTLimit-----number---'
```

Notes:

1. For each Centera device class, you must specify an IP address. However, a Pool Entry Authorization (PEA) file name and path are optional, and the PEA file specification must follow the IP address. Use the "?" character to separate the PEA file name and path from the IP address.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

HLAddress

Specifies an IP address for the Centera storage device and, optionally, the name and path of one Pool Entry Authorization (PEA) file. Specify the IP address with the dotted decimal format (for example, 9.10.111.222). A Centera device might have multiple IP addresses. However, you must specify one of them as a value for this parameter.

AIX The PEA file name and path name are case-sensitive.

If you append the name and path of a PEA file, ensure that the file is stored in a directory on the system that runs the IBM Spectrum Protect™ server. Separate the PEA file name and path from the IP address or addresses with the "?" character, for example: **Windows**

```
HLADDRESS=9.10.111.222?c:\controlFiles\TSM.PEA
```

AIX

```
HLADDRESS=9.10.111.222?/user/ControlFiles/TSM.PEA
```

Specify only one PEA file name and path for each device class definition. If you specify two different Centera device classes that point to the same Centera storage device and if the device class definitions contain different PEA file names and paths, the server uses the PEA file that is specified in the device class HLADDRESS parameter that was first used to open the Centera storage device.

Note:

1. The server does not include a PEA file during installation. If you do not create a PEA file, the server uses the Centera default profile, which can allow applications to read, write, delete, purge, and query data on a Centera storage device. To provide tighter control, create a PEA file with the command-line interface that is provided by EMC Centera. For details about Centera authentication and authorization, refer to the EMC Centera *Programmer's Guide*.
2. You can also specify the PEA file name and path in an environment variable by using the syntax `CENTERA_PEA_LOCATION=filePath_fileName`. The PEA file name and path that is specified with this environment variable apply to all Centera clusters. If you use this variable, you do not need to specify the PEA file name and path using the HLADDRESS parameter.
3. Updating the device class with a new or changed PEA file name and location might require a server restart if the Centera storage device identified by the IP address has already been accessed in the current instance of the server.

MINCAPacity

Specifies the new minimum size for Centera volumes that are assigned to a storage pool in this device class. This value represents the minimum amount of data that is stored on a Centera volume before the server marks it full. Centera volumes

continue to accept data until the minimum amount of data is stored. This parameter is optional.

size

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The minimum value that is allowed is 1 MB (MINCAPACITY=1M). The maximum value that is allowed is 128 GB (MINCAPacity=128G).

MOUNTLimit

Specifies the new maximum number of sessions that access the Centera device. This parameter is optional. You can specify any number from 0 or greater; however, the sum of all mount limit values for all device classes that are assigned to the same Centera device must not exceed the maximum number of sessions that are allowed by Centera.

UPDATE DEVCLASS (Update a DLT device class)

Use the DLT device class when you are using DLT tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-FORMAT-----+DRIVE-----+'
                                     +-DLT1-----+
                                     +-DLT1C-----+
                                     +-DLT10-----+
                                     +-DLT10C-----+
                                     +-DLT15-----+
                                     +-DLT15C-----+
                                     +-DLT20-----+
                                     +-DLT20C-----+
                                     +-DLT35-----+
                                     +-DLT35C-----+
                                     +-DLT40-----+
                                     +-DLT40C-----+
                                     +-DLT2-----+
                                     +-DLT2C-----+
                                     +-DLT4-----+
                                     +-DLT4C-----+
                                     +-SDLT-----+
                                     +-SDLTC-----+
                                     +-SDLT320---+
                                     +-SDLT320C--+
                                     +-SDLT600---+
                                     +-SDLT600C--+
                                     +-DLTS4-----+
                                     '-DLTS4C---'

>--+-----+----->
  '-ESTCAPacity---size-'

>--+-----+----->
  '-PREFIX-----+ADSM-----+'
    '-tape_volume_prefix-'

>--+-----+--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'

>--+-----+-----><
  '-MOUNTLimit-----+DRIVES-+-'
    +-number-+
    '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the DLT tape drives that can be used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

The following table lists the recording formats and estimated capacities for DLT devices:

Table 1. Recording format and default estimated capacity for DLT

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
DLT1	40.0 GB	Uncompressed format, using only CompacTape III or CompacTape IV cartridges
DLT1C	See note 1. 80.0 GB	Compressed format, using only CompacTape III and CompacTape IV cartridges
DLT10	10.0 GB	Uncompressed format, using only CompacTape III or CompacTape IV cartridges
DLT10C	See note 1. 20.0 GB	Compressed format, using only CompacTape III and CompacTape IV cartridges
DLT15	15.0 GB	Uncompressed format, using only CompacTape IIIxt or CompacTape IV cartridges (not CompacTape III) Note: Valid with DLT2000XT, DLT4000, and DLT7000 drives
DLT15C	See note 1. 30.0 GB	Compressed format, using only CompacTape IIIxt or CompacTape IV cartridges (not CompacTape III) Valid with DLT2000XT, DLT4000, and DLT7000 drives
DLT20	20.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT20C	See note 1. 40.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT35	35.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT35C	See note 1. 70.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives

Format	Estimated Capacity	Description
DLT40	40.0 GB	Uncompressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT40C	See note 1. 80.0 GB	Compressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT2	80.0 GB	Uncompressed format, using Quantum DLT tape VS1 media
DLT2C	See note 1. 160.0 GB	Compressed format, using Quantum DLT tape VS1 media
DLT4	160.0 GB	Uncompressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
DLT4C	See note 1. 320.0 GB	Compressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
SDLT See note 2.	100.0 GB	Uncompressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLTC See note 2.	See note 1. 200.0 GB	Compressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLT320 See note 2.	160.0 GB	Uncompressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT320C See note 2.	See note 1. 320.0 GB	Compressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT600	300.0 GB	Uncompressed format, using SuperDLTtape-II media Valid with a Super DLT drive
SDLT600C	See note 1. 600.0 GB	Compressed format, using SuperDLTtape-II media Valid with a Super DLT drive
DLTS4	800 GB	Uncompressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive
DLTS4C	See note 1. 1.6 TB	Compressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive
<p>Note:</p> <ol style="list-style-type: none"> Depending on the effectiveness of compression, the actual capacity might be greater than the listed value. IBM Spectrum Protect™ does not support a library that contains both Backward Read Compatible (BRC) SDLT and Non-Backward Read Compatible (NBRC) SDLT drives. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about estimated capacities, see Table 1.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

UPDATE DEVCLASS (Update an ECARTRIDGE device class)

Use the ECARTRIDGE device class when you are using StorageTek drives such as the StorageTek T9840 or T10000.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see UPDATE DEVCLASS (Update an ECARTRIDGE device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+----->
  '-LIBRARY---library_name-'
>--+-----+----->
  '-LBProtect---+READWrite+-'
                    +-WRITEOnly+
                    '-No-----'
>--+-----+----->
  '-FORMAT---+DRIVE---+' '-ESTCAPacity---size-'
                    +-T9840C---+
                    +-T9840C-C--+
                    +-T9840D---+
                    +-T9840D-C--+
                    +-T10000A---+
                    +-T10000A-C+
                    +-T10000B---+
                    +-T10000B-C+
                    +-T10000C---+
                    +-T10000C-C+
                    +-T10000D---+
                    '-T10000D-C-'
>--+-----+----->
  '-PREFIX---+ADSM-----+'
                    '-tape_volume_prefix-'
>--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>--+-----+----->
  '-MOUNTLimit---+DRIVES--+
                    +-number+
                    '-0-----'
>--+-----+-----><
  | (1) (2) |
  '------DRIVEEncryption---+ON-----+'
                    +-ALLOW-----+
                    +-EXTERNAL+
                    '-OFF-----'
```

Notes:

1. You can use drive encryption only for Oracle StorageTek T10000B drives with a format value of DRIVE, T10000B, or T10000B-C, for Oracle StorageTek T10000C drives with a format value of DRIVE, T10000C or T10000C-C, and for Oracle StorageTek T10000D drives with a format value of DRIVE, T10000D and T10000D-C.
2. You cannot specify both WORM=YES and DRIVEENCRYPTION=ON.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object with the ECARTRIDGE tape drives that can be used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When LBPROTECT is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

READWRITE

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect™ and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the BACKUP DB command.

When the LBPROTECT parameter is set to READWRITE, you do not have to specify the CRCDATA parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEONLY

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the BACKUP DB command.

No

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on Oracle StorageTek T10000C and Oracle StorageTek T10000D drives.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

Important: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format that is established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for ECARTRIDGE devices:

Table 1. Recording formats and default estimated capacities for ECARTRIDGE tapes

Format	Estimated Capacity	Description
--------	--------------------	-------------

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
Notes:		
<ul style="list-style-type: none"> Some formats use a tape drive hardware compression feature. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value. T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats. T10000D drives can read, but cannot write, the T10000A, T10000B, and T10000C formats. 		

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about the default estimated capacity for cartridge tapes, see Table 1.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional.

Restriction:

1. You can use drive encryption only for the following drives:
 - Oracle StorageTek T10000B drives that have a format value of DRIVE, T10000B, or T10000B-C
 - Oracle StorageTek T10000C drives that have a format value of DRIVE, T10000C, or T10000C-C
 - Oracle StorageTek T10000D drives that have a format value of DRIVE, T10000D, or T10000D-C

2. You cannot specify IBM Spectrum Protect as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both WORM=YES and DRIVEENCRYPTION=ON is not supported.)
3. If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, back up sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not allowed and backup operations fail.

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if another method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption. By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable another method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

UPDATE DEVCLASS (Update a FILE device class)

Use the FILE device class when you are using files on magnetic disk storage as volumes that store data sequentially (as on tape).

AIX | **Linux** The FILE device class does not support EXTERNAL libraries.

Windows The FILE device class does not support EXTERNAL libraries.

AIX | **Linux** If you are defining a device class for devices that are to be accessed through a z/OS® media server, see UPDATE DEVCLASS (Update a FILE device class for z/OS media server).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+-----+-----+----->
  '-MOUNTLimit----number-' '-MAXCAPacity----size-'
>--+-----+-----+-----+----->
  |           .-,'-----'. |
  |           v             | |
  '-DIRectory----directory_name-+-'
>--+-----+-----+-----+-----><
  '-SHAREd----+No--+-'
                '-Yes-'
```

Parameters

device_class_name (Required)
Specifies the name of the device class to be updated.

MOUNTLimit

Specifies the maximum number of files that can be simultaneously open for input and output. This parameter is optional. You can specify a number from 0 to 4096.

Windows If the device class is shared with a storage agent (by specifying the SHARED=YES parameter), drives are defined or deleted to match the MOUNTLIMIT value.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

MAXCAPacity

Specifies the maximum size of any data storage files that are categorized by this device class. This parameter is optional.

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The minimum size is 1 MB (MAXCAPACITY=1M). If you are defining a FILE device class for database-backup volumes, specify a value for MAXCAPACITY that is appropriate for the size of the database and that minimizes the number of database volumes.

For example, MAXCAPACITY=5G specifies that the maximum capacity for a volume in this device class is 5 gigabytes. The value that is specified must be less than or equal to the maximum supported size of a file on the target file system.

AIX | **Linux** Do not define a MAXCAPACITY value greater than 640M when this file is for REMOVABLEFILE CD support. A value less than a CD's usable space (650 MB) allows for a one-to-one match between files from the FILE device class and copies that are on CD.

DIRectory

Specifies the directory location or locations of the files that are used in this device class. Enclose the entire list of directories within quotation marks, by using commas to separate individual directory names. Special characters (for example, blank spaces) are allowed within directory names. For example, the directory list "abc def,xyz" contains two directories: abc def and xyz. This parameter is optional.

By specifying a directory name or names, you identify the locations where the server places the files that represent storage volumes for this device class.

AIX | **Linux** While the command is processed, the server expands the specified directory name or names into their fully qualified forms, starting from the root directory.

Important: If you are using storage agents for shared access to FILE volumes, you must use the DEFINE PATH command to define a path for each storage agent. The path definition includes the directory names that are used by the storage agent to access each directory.

Later, if the server must allocate a scratch volume, it creates a new file in one of these directories. (The server can choose any of the directories in which to create new scratch volumes.) For scratch volumes used to store client data, the file that is created by the server has a file name extension of .bfs. For scratch volumes used to store export data, a file name extension of .exp is used.

AIX | **Linux** For example, if you define a device class with a directory of tsmstor and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named /tsmstor/00566497.exp.

Windows For example, if you define a device class with a directory of c:\server and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named c:\server\00566497.exp.

Tip: If you specify multiple directories for a device class, ensure that the directories are associated with separate file systems. Space trigger functions and storage pool space calculations take into account the space that remains in each directory. If you specify multiple directories for a device class and the directories are in the same file system, the server calculates space by adding values that represent the space that remains in each directory. These space calculations are inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the wrong storage pool and run out of space prematurely. For space triggers, an inaccurate calculation might result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled. If a trigger is disabled because the space in a storage pool was not expanded, you can re-enable the trigger by issuing the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

Restriction: To modify a list of directories, you must replace the entire list.

SHARed

Specifies that this FILE device class is shared between the server and one or more storage agents. To prepare for sharing, a library is automatically defined along with a number of drives corresponding to the MOUNTLIMIT associated with the

device class. If the library and drives exist and the MOUNTLIMIT is changed, drives can either be created to reach a new higher MOUNTLIMIT value or deleted to reach a new lower value.

Storage agents using FILE volumes

You must ensure that storage agents can access newly created FILE volumes. To access FILE volumes, storage agents replace names from the directory list in the device-class definition with the names in the directory list for the associated path definition. The following illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created FILE volumes.

Suppose you want to use these three directories for a FILE library:

Windows

- c:\server
- d:\server
- e:\server

AIX

- /usr/tivoli1
- /usr/tivoli2
- /usr/tivoli3

Linux

- /opt/tivoli1
- /opt/tivoli2
- /opt/tivoli3

1. You use the following command to set up a FILE library named CLASSA with one drive named CLASSA1 on SERVER1:

Windows

```
define devclass classa devtype=file
directory="c:\server,d:\server,e:\server"
shared=yes mountlimit=1
```

AIX

```
define devclass classa devtype=file
directory="/usr/tivoli1,/usr/tivoli2,/usr/tivoli3"
shared=yes mountlimit=1
```

Linux

```
define devclass classa devtype=file
directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
shared=yes mountlimit=1
```

2. You want the storage agent STA1 to be able to use the FILE library, so you define the following path for storage agent STA1:

o Windows

```
define path server1 stal srctype=server desttype=drive device=file
directory="\\192.168.1.10\c\server,\\192.168.1.10\d\server,
\\192.168.1.10\e\server" library=classa
```

In this scenario, the storage agent, STA1, replaces the directory name c:\server with the directory name \\192.168.1.10\c\server to access FILE volumes that are in the c:\server directory on the server.

o AIX

```
define path server1 stal srctype=server desttype=drive device=file
directory="/usr/ibm1,/usr/ibm2,/usr/ibm3" library=classa
```

In this scenario, the storage agent, STA1, replaces the directory name /usr/tivoli1 with the directory name /usr/ibm1 to access FILE volumes that are in the /usr/tivoli1 directory on the server.

o Linux

```
define path server1 stal srctype=server desttype=drive device=file
directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

In this scenario, the storage agent, STA1, replaces the directory name /opt/tivoli1 with the directory name /opt/ibm1/ to access FILE volumes that are in the /opt/tivoli1 directory on the server.

The following results occur:

- **Windows** File volume c:\server\file1.dsm is created by SERVER1. If you later change the first directory for the device class with the following command:

```
update devclass classa directory="c:\otherdir,d:\server,e:\server"
```

SERVER1 is still able to access file volume c:\server\file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list that is associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

- **AIX** If file volume /usr/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/usr/otherdir,/usr/tivoli2,
/usr/tivoli3"
```

SERVER1 is still able to access file volume /usr/tivoli1/file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list that is associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

- **Linux** If file volume /opt/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,
/opt/tivoli3"
```

SERVER1 is still able to access file volume /opt/tivoli1/file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list that is associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

Example: Update a FILE device class for sharing

Prepare a FILE device class (named PLAINFILES) for sharing with an IBM Spectrum Protect™ storage agent.

```
update devclass plainfiles shared=yes
```

Example: Update the capacity of a FILE device class

Update a file device class named STORFILES to a maximum capacity of 25 MB.

```
update devclass storfiles maxcap=25m
```

AIX

Example: Add a directory to a FILE device class

Update the FILE device class, CLASSA, by adding a directory, /usr/otherdir, to the directory list. The directories /opt/tivoli2 and /opt/tivoli3 were specified when the device class was first defined.

```
update devclass classa
directory="/opt/tivoli2,/opt/tivoli3,/usr/otherdir"
```

Linux

Example: Add a directory to a FILE device class

Update the FILE device class, CLASSA, by adding a directory, /usr/otherdir, to the directory list. The directories /usr/tivoli2 and /usr/tivoli3 were specified when the device class was first defined.

```
update devclass classa
directory="/usr/tivoli2,/usr/tivoli3,/usr/otherdir"
```

Windows

Example: Add a directory to a FILE device class

Update the FILE device class, CLASSA, by adding a directory, c:\otherdir, to the directory list. The directories d:\server and e:\server were specified when the device class was first defined.

```
update devclass classa
directory="d:\server,e:\server,c:\otherdir"
```

AIX Windows

UPDATE DEVCLASS (Update a GENERICTAPE device class)

Use the GENERICTAPE device class for tape drives that are supported by operating system device drivers.

When this device type is used, the server does not recognize either the type of device or the cartridge recording format. Because the server does not recognize the type of device, if an I/O error occurs, error information is less detailed compared to error information for a specific device type (for example, 8MM). When you define devices to the server, do not mix various types of devices within the same device type.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-ESTCAPacity---size-'
>--+-----+--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>--+-----+-----><
  '-MOUNTLimit---+DRIVES--+
                    +-number+
                    '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the tape drives that can be used by this device class.

This parameter is optional.

For information about defining a library object, see the DEFINE LIBRARY command.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

Specify a capacity appropriate to the particular tape drive that is being used.

You must specify this value as an integer followed by one of the following unit indicators: K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

UPDATE DEVCLASS (Update an LTO device class)

Use the LTO device class when you are using LTO tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+----->
'-LIBRARY----library_name-'
```

```

>----->
'-LBProtect-----+READWrite+-'
              +-WRITEOnly+
              '-No-----'

>----->
|          (1)          | '-ESTCAPacity----size-'
'-FORMAT-----+DRIVE-----'
              +-ULTRIUM2--+
              +-ULTRIUM2C--+
              +-ULTRIUM3--+
              +-ULTRIUM3C--+
              +-ULTRIUM4--+
              +-ULTRIUM4C--+
              +-ULTRIUM5--+
              +-ULTRIUM5C--+
              +-ULTRIUM6--+
              +-ULTRIUM6C--+
              +-ULTRIUM7--+
              +-ULTRIUM7C--+
              +-ULTRIUM8--+
              '-ULTRIUM8C-'

>----->
'-PREFIX-----+ADSM-----'
              '-tape_volume_prefix-'

>----->
'-MOUNTRetention----minutes-' '-MOUNTWait----minutes-'

>----->
'-MOUNTLimit-----+DRIVES+-'
              +-number+
              '-0-----'

>-----><
| (2) (3) |
|-----DRIVEEncryption-----+ON-----|
|-----+-----+
|-----+ALLOW-----+
|-----+EXTERNAL--+
|-----+OFF-----'

```

Notes:

1. IBM Spectrum Protect™ server supports LTO-2 tape drives; however, IBM® Tape Device drivers do not. In the event of an issue with the LTO-2 drive, the preferred corrective action is to upgrade your tape drive hardware to a higher generation drive, then install the latest version of the device driver.
2. You cannot specify DRIVEENCRYPTION=ON if your drives are using WORM (write once, read many) media.
3. Drive encryption is supported only for LTO-4 and higher generation LTO drives and media.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of the defined library object that contains the LTO tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When LBPROTECT is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for IBM Spectrum Protect and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the BACKUP DB command.

When the LBPROTECT parameter is set to READWRITE, you do not have to specify the CRCDATA parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for IBM Spectrum Protect to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the BACKUP DB command.

No

Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction:

Restrictions apply to logical block protection (LBP):

- At the LTO-5 level, LBP is supported only on IBM LTO-5.
- Starting with LTO-6, LBP is supported by all LTO drive vendors.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Instead, specify the format that the drives use.

- If you plan to upgrade all drives to Generation 4, 5, 6, 7, or 8, you must delete all existing LTO Ultrium drive definitions and the paths that are associated with them. Then, you can define the new Generation 4, 5, 6, 7, or 8 drives and paths.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 with LTO-8 drives and media in a single library, you must partition the library into two libraries. One library must have only LTO-8 drives and media and the other must have LTO-6 drives and media.

If you are considering mixing different generations of LTO media and drives, be aware of the following restrictions.

Table 1. Read - write capabilities for different generations of LTO drives

Drives	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media	Generation 7 media	Generation M8 media	Generation 8 media
Generation 3 ¹	Read and write	n/a	n/a	n/a	n/a	n/a	n/a
Generation 4 ¹	Read and write	Read and write	n/a	n/a	n/a	n/a	n/a
Generation 5 ¹	Read only	Read and write	Read and write	n/a	n/a	n/a	n/a
Generation 6 ¹	n/a	Read only	Read and write	Read and write	n/a	n/a	n/a
Generation 7 ¹			Read only	Read and write	Read and write	n/a	n/a
Generation 8 ²	n/a	n/a	n/a	n/a	Read and write	Read and write	Read and write

Drives	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media	Generation 7 media	Generation M8 media	Generation 8 media
¹ If a storage pool volume can only be read by a tape drive, ensure that the attributes of the storage pool volume are set to read only.							
² LTO-8 drives have two media types: LTO-M8 media and LTO-8 media. Both media types are used only in LTO-8 tape drives.							

The following table lists the recording formats and estimated capacities for LTO devices:

Table 2. Recording format and default estimated capacity for LTO

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
ULTRIUM2	200 GB	Uncompressed (standard) format, using Ultrium 2 cartridges
ULTRIUM2C	See note 400 GB	Compressed format, using Ultrium 2 cartridges
ULTRIUM3	400 GB	Uncompressed (standard) format, using Ultrium 3 cartridges
ULTRIUM3C	See note 800 GB	Compressed format, using Ultrium 3 cartridges
ULTRIUM4	800 GB	Uncompressed (standard) format, using Ultrium 4 cartridges
ULTRIUM4C	See note 1.6 TB	Compressed format, using Ultrium 4 cartridges
ULTRIUM5	1.5 TB	Uncompressed (standard) format, using Ultrium 5 cartridges
ULTRIUM5C	Varied, as described in note	Compressed format, using Ultrium 5 cartridges
ULTRIUM6	2.5 TB	Uncompressed (standard) format, using Ultrium 6 cartridges
ULTRIUM6C	Varied, as described in note	Compressed format, using Ultrium 6 cartridges
ULTRIUM7	6 TB	Uncompressed (standard) format, using Ultrium 7 cartridges
ULTRIUM7C	Varied, as described in note	Compressed format, using Ultrium 7 cartridges
ULTRIUM8	12 TB for LTO-8 media 9 TB for LTO-M8 media	Uncompressed (standard) format, using Ultrium M8 or Ultrium 8 cartridges
ULTRIUM8C	Varied, as described in note	Compressed format, using Ultrium M8 or Ultrium 8 cartridges
Note: If this format uses the tape-drive hardware-compression feature, depending on the effectiveness of compression, the actual capacity is varied.		

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes that are categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about estimated capacities, see Table 2.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

DRIVEEncryption

Specifies whether drive encryption is allowed. This parameter is optional. Drive encryption is supported only for LTO-4 and higher generation drives and media.

Restriction: If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON

Specifies that IBM Spectrum Protect is the key manager for drive encryption and allows drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, back up sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not allowed and backup operations fail.

Note: You cannot specify IBM Spectrum Protect as the key manager for drive encryption of WORM (write once, read many) media. (If you are using WORM media, you cannot specify DRIVEENCRYPTION=ON.)

ALLOW

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. However, drive encryption for empty volumes is allowed if another method of encryption is enabled.

EXTERNAL

Specifies that IBM Spectrum Protect does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect does not turn off encryption. By contrast, when you specify ALLOW and IBM Spectrum Protect detects that AME encryption is enabled, IBM Spectrum Protect turns off encryption.

OFF

Specifies that drive encryption is not allowed. If you enable another method of encryption, backups fail. If you enable the application method, IBM Spectrum Protect disables encryption and backups are attempted.

Example: Update the mount limit for an LTO device class

Update a device class named LTOTAPE. Change the mount limit to 2.

```
update devclass ltotape mountlimit=2
```

UPDATE DEVCLASS (Update a NAS device class)

Use the NAS device class when you are using NDMP (Network Data Management Protocol) operations to back up network-attached storage (NAS) file servers. The device class is for drives that are supported by the NAS file server for backups.

AIX | **Linux** The NAS device class does not support EXTERNAL libraries.

Windows The NAS device class does not support EXTERNAL libraries.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-MOUNTRetention---0-'
>--+-----+--+-----+----->
  '-MOUNTWait---minutes-' '-MOUNTLimit---+DRIVES+-'
                               +-number+
                               '-0-----'
```

```

>-----+----->
'-ESTCAPacity-----size-'
>-----+-----<
'-PREFIX-----tape_volume_prefix-'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of the defined library object that contains the SCSI tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

MOUNTRetention=0

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. Zero (0) is the only supported value for device classes with DEVType=NAS.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

Example: Update the estimated capacity for a NAS device class

Update a device class named NASTAPE. Change the estimated capacity to 200 GB.

```
update devclass nastape library=naslib estcapacity=200G
```

UPDATE DEVCLASS (Update a REMOVABLEFILE device class)

Use the REMOVABLEFILE device class for removable media devices that are attached as local, removable file systems.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-MAXCAPacity---size-'
>--+-----+--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>--+-----+-----><
  '-MOUNTLimit---+DRIVES-+-'
                    +-number-+
                    '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the removable media drives used by this device class. This parameter is optional. For information about defining a library object, see the DEFINE LIBRARY command.

MAXCAPacity

Specifies the maximum size of any volumes that are defined to a storage pool categorized by this device class. This parameter is optional.

AIX | **Windows** Because the server opens only one file per physical removable medium, specify a capacity that enables one file to make full use of your media capacity.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes).

For example, MAXCAPACITY=5M specifies that the maximum capacity for a volume in this device class is 5 MB. The smallest value that is allowed is 1 MB (that is, MAXCAPACITY=1M).

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

UPDATE DEVCLASS (Update a SERVER device class)

Use the SERVER device class to use storage volumes or files that are archived in another IBM Spectrum Protect™ server.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDdate DEVclass--device_class_name----->
>--+-----+-----+-----+----->
  '-SERVERName----server_name-' '-MAXCAPacity----size-'
>--+-----+-----+-----+----->
  '-PREFIX----+ADSM-----+-'
                   '-tape_volume_prefix-'
>--+-----+-----+-----+----->
  '-RETRYPeriod-----minutes--'
```

```

>----->
'-RETRYInterval-----seconds--'
>----->
'-MOUNTRetention-----minutes-'
>----->>
'-MOUNTLimit-----+number-+-'
'-1-----'

```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

SERVERName

Specifies the name of the server. The SERVERNAME parameter must match a defined server.

Note: If you change the SERVERNAME of an existing server to a new name, data on the volumes under the old SERVERNAME is no longer accessible with this device class.

MAXCAPacity

Specifies the maximum size that objects can be when created on the target server. This parameter is optional.

Specify this value as an integer followed by K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). The minimum value that is allowed is 1 MB (MAXCAPACITY=1M).

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

```
AB.CD2.E
```

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high-level archive file name that uses the default prefix is ADSM.volume1.

RETRYPeriod

Specifies the retry period in minutes. The retry period is the interval during which the server attempts to contact a target server if there is a suspected communications failure. This parameter is optional. You can specify a number 0 - 9999.

RETRYInterval

Specifies the retry interval in seconds. The retry interval is how often retries are done within a specific time period. This parameter is optional. You can specify a number 1 - 9999.

MOUNTRetention

Specifies the number of minutes to retain an idle connection with the target server before the connection is closed. This parameter is optional. You can specify a number 0 - 9999.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTLimit

Specifies the maximum number of simultaneous sessions between the source server and the target server. Any attempts to access more sessions than indicated by the mount limit cause the requester to wait. This parameter is optional. You can specify a number 1 - 4096.

The following are possible values:

number

- 1 Specifies the maximum number of simultaneous sessions between the source server and the target server.
- Specifies the number of simultaneous sessions between the source server and the target server.

UPDATE DEVCLASS (Update a VOLSAFE device class)

Use the VOLSAFE device type to work with StorageTek VolSafe brand media and drives. This technology uses media that cannot be overwritten. Therefore, do not use these media for short-term backups of client files, the server database, or export tapes.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+-----+-----+----->
  '-LIBRARY----library_name-'  '-FORMAT-----+DRIVE-----+'
                                     +-9840-----+
                                     +-9840-C----+
                                     +-T9840C----+
                                     +-T9840C-C--+
                                     +-T9840D----+
                                     +-T9840D-C--+
                                     +-T10000A---+
                                     +-T10000A-C+
                                     +-T10000B---+
                                     +-T10000B-C+
                                     +-T10000C---+
                                     +-T10000C-C+
                                     +-T10000D---+
                                     +-T10000D-C-'
>--+-----+-----+-----+----->
  '-ESTCAPacity----size-'
>--+-----+-----+-----+----->
  '-PREFIX-----+ADSM-----+-'
                    '-tape_volume_prefix-'
>--+-----+-----+-----+----->
  '-MOUNTRetention---minutes-'  '-MOUNTWait----minutes-'
>--+-----+-----+-----+-----><
  '-MOUNTLimit-----+DRIVES--+-'
                        +-number-+
                        '-0-----'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of the defined library object that contains the VolSafe drives that can be used by this device class. If any drives in a library are VolSafe-enabled, all drives in the library must be VolSafe-enabled. For more information about the VolSafe device type, see DEFINE DEVCLASS (Define a VOLSAFE device class).

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

Attention: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format that is established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for VolSafe devices:

Table 1. Recording formats and default estimated capacities for volsafe tapes

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
9840	20 GB	Uncompressed (standard) format, using a 20 GB cartridge with 270 meters (885 feet) of tape
9840-C	80 GB	LZ-1 Enhanced (4:1) compressed format, using an 80 GB cartridge with 270 meters (885 feet) of tape
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by one of the following unit indicators: **K** (kilobytes), **M** (megabytes), **G** (gigabytes), or **T** (terabytes). The smallest value that is accepted is 1 MB (ESTCAPACITY=1M).

For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G.

To force the IBM Spectrum Protect™ server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about the default estimated capacity for cartridge tapes, see Table 1.

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high-level archive file name that uses the default prefix is ADSM.volume1.

MOUNTRetention

Specifies the number of minutes that an idle sequential access volume is retained before it is dismounted. This parameter is optional. You can specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types (that is, a library that is managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

Note: For environments in which devices are shared across storage applications, the MOUNTRETENTION setting must be carefully considered. This parameter determines how long an idle volume remains in a drive. Some media managers do not dismount an allocated drive to satisfy pending requests. You might need to tune this parameter to satisfy competing mount requests while maintaining optimal system performance. Typically, problems arise more frequently when the MOUNTRETENTION parameter is set to a value that is too small, for example, zero.

MOUNTWait

Specifies the maximum number of minutes the server waits for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number 0 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transactions continue and complete, but new transactions are terminated.

AIX | Linux

UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server)

Use this command to update a device class. A limited set of device class types is available for devices that are accessed through a z/OS® media server.

- UPDATE DEVCLASS (Update a 3590 device class for z/OS media server)

- UPDATE DEVCLASS (Update a 3592 device class for z/OS media server)
- UPDATE DEVCLASS (Update an ECARTRIDGE device class for z/OS media server)
- UPDATE DEVCLASS (Update a FILE device class for z/OS media server)

Table 1. Commands related to UPDATE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up IBM Spectrum Protect device information to a file.
DEFINE DEVCLASS (z/OS media server)	Defines a device class to use storage managed by a z/OS media server.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
UPDATE LIBRARY	Changes the attributes of a library.

AIX Linux

UPDATE DEVCLASS (Update a 3590 device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS® media server to access 3590 devices. The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```

(1) (2)
>>-UPDate DEVclass--device_class_name----->
>--+-----+--+-----+----->
  '-LIBRARY---library_name-' '-FORMAT---+DRIVE---+'
                                     +-3590B---+
                                     +-3590C---+
                                     +-3590E-B++
                                     +-3590E-C++
                                     +-3590H-B++
                                     '-3590H-C-'
>--+-----+--+-----+----->
  '-ESTCAPacity---size-' '-COMPRESSION---+Yes-+-'
                                     '-No--'
>--+-----+--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>--+-----+--+-----+----->
  '-MOUNTLimit---+DRIVES-+-' '-EXPIRATION---yyyyddd-'
                                     +-number-+
                                     '-0-----'
>--+-----+--+-----+----->
  '-RETENTION---days-' '-PROTECTION---+No-----+-'
                                     +-Yes-----+
                                     '-Automatic-'
>--+-----+-----><
  '-UNIT---unit_name-'

```

Notes:

1. You must specify at least one optional parameter on this command.

- You cannot update the PREFIX parameter with this command. You must create a device class with the value that you require for the PREFIX parameter.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

This parameter is optional.

For information about defining a library, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The following table lists the recording format options for 3590 devices:

Table 1. Recording formats for 3590

Format	Description
3590B	Uncompressed (basic) format
3590C	Compressed format
3590E-B	Uncompressed (basic) format, similar to the 3590B format
3590E-C	Compressed format, similar to the 3590C format
3590H-B	Uncompressed (basic) format, similar to the 3590B format
3590H-C	Compressed format, similar to the 3590C format

Note: If the format uses the tape drive hardware compression feature the actual capacity can increase, depending on the effectiveness of compression.

ESTCAPACITY

Specifies the estimated capacity for the sequential access volumes that are categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), G (GB), or T (TB). For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G. The smallest value that is accepted is 100 KB (ESTCAPACITY=100K).

COMPRESSION

Specifies whether file compression is used for this device class. This parameter is optional.

You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

MOUNTRETENTION

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWAIT

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyyddd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

Tip: You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the EXPIRATION parameter. You cannot specify a value for the EXPIRATION parameter if you specify a non-zero value for the RETENTION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive,

allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3590 tape. This parameter is optional. The unit name can be up to 8 characters.

AIX | Linux

UPDATE DEVCLASS (Update a 3592 device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS® media server to access 3592 devices. The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```

                                     (1) (2)
>>-UPDate DEVclass--device_class_name----->
>--+-----+----->
  '-LIBRARY-----zos_media_library-'
>--+-----+----->
  '-FORMAT-----DRIVE---+' '-ESTCAPacity-----size-'
      +-3592-----
      +-3592C---+
      +-3592-2---+
      +-3592-2C--+
      +-3592-3---+
      +-3592-3C--+
      +-3592-4---+
      '-3592-4C-'
>--+-----+----->
```

```

'-COMPression-----+Yes+-'
                    '-No--'
>-----+-----+-----+----->
'-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>-----+-----+-----+----->
'-MOUNTLimit-----+DRIVES+-' '-EXPIration---yyyddd-'
                    +-number-+
                    '-0-----'
>-----+-----+-----+----->
'-RETention---days-' '-PROtection---+No-----+-'
                                +-Yes-----+
                                '-Automatic-'
>-----+-----+-----+----->>
'-UNIT---unit_name-'

```

Notes:

1. You must specify at least one optional parameter on this command.
2. You cannot update the PREFIX parameter with this command. You must create a device class with the value that you require for the PREFIX parameter.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

This parameter is optional.

For information about defining a library, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. The default value is DRIVE.

See the following table for the recording formats.

Table 1. Recording formats for 3592

Format	Description
3592	Uncompressed (basic) format
3592C	Compressed format
3592-2	Uncompressed (basic) format, similar to the 3592 format
3592-C	Compressed format, similar to the 3592C format
3592-3	Uncompressed (basic) format, similar to the 3592 format
3592-3C	Compressed format, similar to the 3592C format
3592-4	Uncompressed (basic) format, similar to the 3592 format
3592-4C	Compressed format, similar to the 3592C format
DRIVE	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be different from the listed value.	

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. For optimal results, do not mix generations of drives in the same library. If a library contains mixed generations, media problems can result. For example, generation 1 and generation 2 drives cannot read generation 3 media. If possible, upgrade all drives to 3592 generation 3. If you cannot upgrade all drives to 3592 generation 3, you must use a special configuration.

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: **K** (KB), **M** (MB), **G** (GB), or **T** (TB). For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G. The smallest value that is accepted is 100 KB (ESTCAPACITY=100K).

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

EXpiration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyyddd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

Tip: You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the EXPIRATION parameter. You cannot specify a value for the EXPIRATION parameter if you specify a non-zero value for the RETENTION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3592 tape. This parameter is optional. This name can be as many as 8 characters.

AIX

Linux

UPDATE DEVCLASS (Update an ECARTRIDGE device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS® media server to access StorageTek drives such as the StorageTek T9840 or T10000. The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
(1) (2)
>>-UPDate DEVclass--device_class_name----->
>--+-----+----->
  '-LIBRARY---zos_media_library-'
>--+-----+----->
  '-FORMAT---+DRIVE---+' '-ESTCAPacity---size-'
      +-T9840C---+
      +-T9840C-C--+
      +-T9840D---+
      +-T9840D-C--+
      +-T10000A---+
      +-T10000A-C+
      +-T10000B---+
      +-T10000B-C+
      +-T10000C---+
      +-T10000C-C+
      +-T10000D---+
      '-T10000D-C-'
>--+-----+----->
  '-MOUNTRetention---minutes-' '-MOUNTWait---minutes-'
>--+-----+----->
  '-MOUNTLimit---+DRIVES--+-' '-COMPRESSION---+Yes--+-'
      +-number+
      '-0-----'
      '-No--'
>--+-----+----->
  '-EXPIration---yyyddd-' '-RETention---days-'
>--+-----+-----><
  '-PROtection---+No-----+' '-UNIT---unit_name-'
      +-Yes-----+
      '-Automatic-'
```

Notes:

1. You must specify at least one optional parameter on this command.
2. You cannot update the PREFIX parameter with this command. You must create a device class with the value that you require for the PREFIX parameter.

Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of a library that was defined with the LIBTYPE=ZOSMEDIA parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

This parameter is optional.

For information about defining a library, see the DEFINE LIBRARY command.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional. See the following table for the recording formats.

Table 1. Recording formats for ECARTRIDGE tapes

Format	Estimated Capacity	Description
DRIVE	-	The server selects the highest format that is supported by the drive on which a volume is mounted. DRIVE is the default value. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000D	8 TB	Uncompressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
T10000D-C	15 TB	Compressed T10000D format, using an Oracle StorageTek T10000 T2 cartridge
Note:		
<ul style="list-style-type: none">• Some formats use a compression feature of the tape drive hardware. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value.• T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats. T10000D drives can read, but cannot write, the T10000A, T10000B, and T10000C formats.		

ESTCAPACITY

Specifies the estimated capacity for the sequential access volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), G (GB), or T (TB). For example, specify that the estimated capacity is 9 GB with the parameter ESTCAPACITY=9G. The smallest value that is accepted is 100 KB (ESTCAPACITY=100K).

MOUNTRETENTION

Specifies the number of minutes that an idle tape volume is retained before it is dismantled. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

Restriction: If the library that is associated with this device class is external (LIBTYPE=EXTERNAL), do not specify the MOUNTWAIT parameter.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

You can specify one of the following values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional.

You can specify one of the following values:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the EXPIRATION parameter, you cannot specify the RETENTION parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems.

If you specify the RETENTION parameter, you cannot specify the EXPIRATION parameter.

Tip: You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the EXPIRATION parameter. You cannot specify a value for the EXPIRATION parameter if you specify a non-zero value for the RETENTION parameter.

PROtection

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. You can specify one of the following values:

No

Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use PROTECTION=YES and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using PROTECT=YES in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify PROTECTION=AUTOMATIC, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify PROTECTION=AUTOMATIC, the z/OS media server issues RACROUTE commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to PROTECTION=AUTOMATIC for a device class that was set to PROTECTION=NO. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for PROTECTION is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to AUTOMATIC, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support ECARTRIDGE tapes. Use the unit name that represents the subset of drives in the library that are attached to the z/OS system. This parameter is optional. The unit name can be up to 8 characters.

AIX | Linux

UPDATE DEVCLASS (Update a FILE device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS® media server to access files on magnetic disk storage as sequential-access volumes (like tape). The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

A volume in this device class is a Virtual Storage Access Method (VSAM) linear data set that is accessed by the z/OS media server. SCRATCH volumes can be used with a device class and the z/OS media server dynamically allocates the VSAM LDS. It is not necessary to define volumes for the server to use the device class. If you define volumes, set the high-level qualifier (HLQ) so that SMS recognizes the allocation request by the z/OS media server. If you are using defined volumes, the format volume function is

not supported for the server when you use this device class. The z/OS media server z/OS media server uses a FormatWrite feature of DFSMS Media Manager when filling FILE volumes.

You can define volumes for the FILE device class by using the DEFINE VOLUME command. However, the z/OS media server does not allocate space for a defined volume until the volume is opened for its first use.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DEVclass--device_class_name----->
>--+-----+-----+-----+----->
  '-MAXCAPacity---size-' '-PRIMARYalloc---size-'
>--+-----+-----+-----+----->
  '-SECONDARYalloc---size-'
>--+-----+-----+-----+----->
  '-PREFIX---file_volume_prefix-'
>--+-----+-----+-----+----->>
  '-MOUNTLimit---number-'
```

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

MAXCAPacity

Specifies the maximum size of file volumes that are defined to a storage pool in this device class. This parameter is optional.

Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum size is 1 MB (MAXCAPACITY=1M). The maximum size is 16384 GB (MAXCAPACITY=16384G).

PRIMARYalloc

Specifies the initial amount of space that is dynamically allocated when a new volume is opened. Enough space must be available to satisfy the primary allocation amount. Storage Management Subsystem (SMS) policy determines whether multiple physical volumes can be used to satisfy the primary allocation request.

This parameter is optional. Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum size is 100 KB (PRIMARYALLOC=100K). The maximum size is 16384 GB (MAXCAPACITY=16384G). All values are rounded to the next higher multiple of 256 KB.

To avoid wasted space, the dynamic allocation operation uses the smaller of the values that are specified in the two parameters, PRIMARYALLOC and MAXCAPACITY.

SMS automatic class selection (ACS) routines can affect whether the PRIMARYALLOC and SECONDARYALLOC parameter values are used.

SECONDARYalloc

Specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up. The data set for a file volume is extended up to the size set by the MAXCAPACITY parameter, then the volume is marked full.

Because secondary allocation of a linear data set cannot span a physical volume, consider the size of the physical volume when selecting a secondary allocation size. For example, physical volumes for a 3390 Model 3 are approximately 2.8 GB. To ensure that each extend request occupies nearly an entire physical volume but not more, use a secondary allocation size that is just less than 2.8 GB. A secondary allocation amount of 2600 MB allots enough space for the VSAM volume data set (VVDS), the volume label, and the volume table of contents (VTOC).

This parameter is optional. Specify this value as an integer followed by **K** (KB), **M** (MB), **G** (GB), or **T** (TB). The minimum value is 0 KB (SECONDARYALLOC=0K). The maximum value is 16384 GB. Except for 0, all values are rounded to the next higher multiple of 256 KB.

If you specify 0 (SECONDARYALLOC=0), the file volume cannot be extended beyond the primary allocation amount.

SMS automatic class selection (ACS) routines can affect whether the PRIMARYALLOC and SECONDARYALLOC parameter values are used.

If you specify a value for the SECONDARYALLOCATION parameter that is not 0, or if you allow the value to default to 2600M, the SMS DATACLAS associated with the PREFIX identifier (for example, High Level Qualifier) must have the Extended Addressability (EA) attribute specified. Without the EA attribute, the SMS DATACLAS limits the allocation of the VSAM LDS FILE volume to the primary extent. (See the description of the PRIMARYALLOCATION parameter). With the data set limited to primary allocation size, the data set cannot be extended by the z/OS media server, and the volume is marked FULL before the maximum capacity is reached.

Restriction: Ensure that the values that you specify for the PRIMARYALLOC and SECONDARYALLOC parameters are within practical limits for the storage device. The server cannot check whether the values exceed practical device limits, and does not check whether the two values together exceed the current MAXCAPACITY setting.

Tip: To fill volumes when you specify a large value for the MAXCAPACITY parameter, specify large values for the PRIMARYALLOC and SECONDARYALLOC parameters. Use larger MVS™ volume sizes to reduce the chance of extend failure.

PREFIX

Specifies the high-level qualifier of the data set name that is used to allocate scratch volume data sets. For all scratch file volumes created in this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of the prefix, including periods, is 32 characters.

Values that are specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

`AB.CD2.E`

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a file volume data set name using the default prefix is `ADSM.B0000021.BFS`.

If you have a data set naming convention, use a prefix that conforms to your naming conventions. For example, the following value is acceptable: `TSM.SERVER2.VSAMFILE`.

If you are running multiple server instances for either IBM Spectrum Protect™ or Tivoli® Storage Manager for z/OS Media you must use a unique value for the PREFIX parameter for each device class that you update.

MOUNTLimit

Specifies the maximum number of FILE volumes that can be open concurrently for this device class. This parameter is optional. For 3995 devices emulating 3390 devices, the value must not be set higher than the numbers of concurrent input and output streams possible on the media storing the volumes.

The value that you specify in this parameter is important if there is a significant penalty switching from one volume to another. For example, switching can take place when using IBM® 3995 devices to emulate 3390 devices. The value that you specify must be no higher than the number of physical drives available on the device.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the MOUNTLIMIT parameter for a device class, the transaction fails.

UPDATE DOMAIN (Update a policy domain)

Use this command to change a policy domain.

Privilege class

Update the policy domain ENGPOLDOM so that the backup retention grace period is extended to 90 days and the archive retention grace period is extended to two years. Specify an active-data pool as the destination for active versions of backup data belonging to nodes that are assigned to the domain. Use *engactivedata* as the name of the active-data pool. Issue the following command:

```
update domain engpoldom description='Engineering Policy Domain'
backretention=90 archretention=730 activedestination=engactivedata
```

Related commands

Table 1. Commands related to UPDATE DOMAIN

Command	Description
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.
QUERY DOMAIN	Displays information about policy domains.

UPDATE DRIVE (Update a drive)

Use this command to update a drive.

Privilege class

For detailed and current drive support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate DRive--library_name--drive_name----->
>--+-----+-----+-----+----->
  '-Serial---+--serial_number--+-'  '-ONLine---+--Yes--+-'
        '-AUTODetect----'          '-No--'
>--+-----+-----+----->
  '-ELEMent---+--address----+-'
        '-AUTODetect-'
>--+-----+-----+----->
  |                               (1) |
  '-ACSDRVID---+--drive_id-----'
>--+-----+-----+----->>
  |                               (2) |
  '-CLEANFREquency---+--NONE-----+-'
                    |               (3) |
                    +-ASNEEDED-----+
                    '-gigabytes-----'
```

Notes:

1. The ACSDRVID parameter is valid only for drives in ACSLS libraries.
2. The CLEANFREQUENCY parameter is valid only for drives in SCSI libraries.
3. The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. For more information, see the parameter description.

Parameters

library_name (Required)

Specifies the name of the library to which the drive is assigned.

drive_name (Required)

Specifies the name that is assigned to the drive.

SERial

Specifies the serial number for the drives that are being updated. This parameter is valid only for drives in a SCSI or virtual tape library (VTL). This parameter is optional. The possible values are:

serial_number

Specifies the serial number for the drive that is being updated.

Note: If a path to this drive is already defined, then the number you enter here is compared to the number detected by IBM Spectrum Protect™. If the numbers do not match, the command fails.

AUTODETECT

Specifies that the serial number is automatically detected and used by IBM Spectrum Protect if a path is already defined to this drive.

If a path to this drive is not defined, then the serial number is not detected.

ONLine

Specifies whether the drive is available for use. This parameter specifies whether drives can be taken offline and used for another activity, such as maintenance. This parameter is optional.

You can issue this command when the drive is involved in an active process or session, but it is not advised. If you issue a command to take the drive offline while it is in use, an error message is issued. The mounted volume completes its current process. If this volume was part of a series of volumes for a specific transaction, the drive is not available to complete mounting the series. If no other drives are available, the process fails.

Attention: When a drive is in use, do not specify the ELEMENT parameter with the ONLINE parameter. The drive is not updated, and the command fails.

The drive state is not changed even if the server is halted and restarted. If a drive is offline when the server is restarted, a warning message is issued stating that the drive must be manually brought online. If all of the drives in a library are updated to be offline, processes that need a library mount point fail, rather than queue up for a mount point.

YES

Specifies that the drive is available for use (online).

No

Specifies that the drive is not available for use (offline).

ELEMent

Specifies the element address of the drive within a SCSI or VTL library. The server uses the element address to connect the physical location of the drive to the SCSI address of the drive. This parameter is valid only for a drive in a SCSI or VTL library when the command is issued from an IBM Spectrum Protect library manager server. The possible values are:

address

Specifies the element address for the drive that is being updated.

To find the element address for your library configuration, consult the information from the manufacturer.

Remember: If a path to this drive is already defined, then the number you enter here is compared to the number previously detected by IBM Spectrum Protect. If the numbers do not match, then this command fails.

AUTODETECT

Specifies that the element number is automatically detected and used by IBM Spectrum Protect if a path is already defined to this drive.

If a path to this drive is not defined, then the element number is not detected.

Restriction: If the library in which the drive is located does not support the Read Element Status SCSI command, and ELEMENT=AUTODETECT, the command fails with an IBM Spectrum Protect error message.

ACSDRVID

Specifies the ID of the drive that is being accessed in an ACSLS library. The drive ID is a set of numbers that indicates the physical location of a drive within an ACSLS library. This drive ID must be specified as *a,l,p,d*, where *a* is the ACSID, *l* is the

LSM (library storage module), *p* is the panel number, and *d* is the drive ID. The server needs the drive ID to connect the physical location of the drive to the drive's SCSI address. See your StorageTek documentation for details.

CLEANFREQUENCY

Specifies how often the server activates drive cleaning. This parameter is optional. For the most complete automation of cleaning for an automated library, you must have a cleaner cartridge checked into the volume inventory for the library. If you are using library based cleaning, NONE is advised when your library type supports this function. This parameter is valid only for drives in SCSI libraries, and not valid for externally managed libraries, such as 3494 libraries or StorageTek libraries that are managed under ACSLS.

Important: There are special considerations if you plan to use server-activated drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

NONE

Specifies that the server does not track cleaning for this drive. Use this parameter for libraries that have their own automatic cleaning.

ASNEEDED

Specifies that the server loads the drive with a checked-in cleaner cartridge only when a drive reports to the device driver that it needs cleaning.

The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. Visit the Supported Devices website for your operating system to view detailed drive information. If ASNEEDED is not supported, you can use the gigabytes value for automatic cleaning.

For IBM 3592 and LTO drives, library based cleaning is advised. If library based cleaning is not supported, then ASNEEDED must be used. Gigabytes is not recommended.

Restriction: IBM Spectrum Protect does not control the drives that are connected to the NAS file server. If a drive is attached only to a NAS file server (no connection to a storage agent or server), do not specify ASNEEDED for the cleaning frequency.

gigabytes

Specifies, in gigabytes, how much data is processed on the drive before the server loads the drive with a cleaner cartridge. The server resets the gigabytes-processed counter each time it loads a cleaner cartridge in the drive.

Important: When CLEANFREQUENCY=gigabyte, drive cleaning can occur before the gigabyte setting is reached, if the drive notifies the device driver that a cleaning is necessary.

Consult the information from the drive manufacturer for cleaning recommendations. If the information gives recommendations for cleaning frequency in terms of hours of use, convert to a gigabytes value by doing the following:

1. Use the bytes-per-second rating for the drive to determine a gigabytes-per-hour value.
2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

Tip: For IBM 3590, specify a value for the cleaning frequency to ensure that the drives receive adequate cleaning. Consult the information from the drive manufacturer for cleaning recommendations. Using the cleaning frequency that is recommended by IBM does not over clean the drives.

Example: Update the element address for a drive

Update DRIVE3, in the library named AUTO, by changing the element address to 119.

```
update drive auto drive3 element=119
```

Example: Take a drive offline

Update DRIVE3, in the library named MANLIB, to take it offline.

```
update drive manlib drive3 online=no
```

Related commands

Table 1. Commands related to UPDATE DRIVE

Command	Description
CLEAN DRIVE	Marks a drive for cleaning.

1. You cannot specify a file space identifier (FSID) if you use wildcard characters for the client node name.
2. You can specify each rule only once.
3. You must specify either the REPLRULE or the REPLSTATE parameter on this command.
4. The ACTIVE_DATA and ACTIVE_DATA_HIGH_PRIORITY rules are valid only if you specify DATATYPE=BACKUP.

Parameters

node_name (Required)

Specifies the client node to which the file space belongs. You can use wildcard characters to specify this name. However, file space identifiers can be different among client nodes for the same file space. Therefore, you cannot specify wildcard characters for the client node name and FSID as the value for the NAMETYPE parameter.

file_space_name (Required)

Specifies the name of the file space to be updated. You can use wildcard characters or a comma-delineated list to specify names.

For a server that has clients with Unicode-enabled file spaces, you might have to make the server convert the file space name that you enter. For example, you might have to make the server convert a name from the server code page to Unicode. For details, see the NAMETYPE parameter. If you specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

File space names are case-sensitive. To determine the correct capitalization for the file space to be updated, use the QUERY FILESPACE command.

NAMETYPE

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for IBM Spectrum Protect™ clients that Unicode-enabled and that have Windows, Macintosh OS X, or NetWare operating systems. Use this parameter only when you enter a partly-qualified or fully-qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret file space names.

UNICODE

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the operating system, on the characters in the name, and the server code page. Conversion can fail if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion fails, the name can contain question marks, blanks, or ellipses (...).

FSID

The server interprets file space names as file space identifiers.

CODETYPE

Specifies the type of file spaces to be included in node replication processing. The default value is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Specifies only file spaces that are in Unicode.

NONUNICODE

Specifies only file spaces that are not in Unicode.

BOTH

Specifies all file spaces regardless of code page type.

DATATYPE (Required)

Specifies the data type to which a replication rule applies. To specify multiple data types, separate the names with commas and no intervening spaces. You can specify the following values:

BACKUP

Specifies the backup data type.

ARCHIVE

Specifies the archive data type.

SPACEMANAGED

Specifies the space-managed data type.

REPLRULE

Specifies the replication rule that applies to a data type. You cannot use wildcards. If you specify multiple data types, the replication rule applies to each data type. For example, if you specify `DATATYPE=BACKUP, ARCHIVE`, the replication rule applies to backup data and to archive data.

Restriction: The `REPLRULE` parameter is optional. However, if you do not specify it, you must specify the `REPLSTATE` parameter.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that a file space contains active backup data and archive data. Replication of the active backup data is a higher priority than the archive data. To prioritize the active backup data, specify `DATATYPE=BACKUP REPLRULE=ACTIVE_DATA_HIGH_PRIORITY`. To assign a normal priority to archive data, issue the `UPDATE FILESPACE` command again, and specify `DATATYPE=ARCHIVE REPLRULE=ALL_DATA`.

You can specify the following rules:

`ALL_DATA`

Replicates backup, archive, or space-managed data. The data is replicated with a normal priority.

`ACTIVE_DATA`

Replicates only the active backup data in a file space. The data is replicated with a normal priority.

Attention: If you specify `ACTIVE_DATA` and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a server version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the `REPLICATE NODE` command with the `FORCERECONCILE=YES` parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a server version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

`ALL_DATA_HIGH_PRIORITY`

Replicates backup, archive, or space-managed data. The data is replicated with a high priority.

`ACTIVE_DATA_HIGH_PRIORITY`

This rule is the same as the `ACTIVE_DATA` replication rule except data is replicated with a high priority.

`DEFAULT`

Data is replicated according to the client node rule for the data type.

For example, suppose that you want to replicate the archive data in all the file spaces that belong to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify `DATATYPE=ARCHIVE REPLRULE=DEFAULT` for each file space. Ensure that the client replication rule for archive data is set to `ALL_DATA_HIGH_PRIORITY` or to `DEFAULT`. If the client replication rule is `DEFAULT`, the server replication rule for archive data must be set to `ALL_DATA_HIGH_PRIORITY`.

`NONE`

Data is not replicated. For example, if you do not want to replicate the space-managed data in a file space, specify `DATATYPE=SPACEMANAGED REPLRULE=NONE`.

`REPLState`

Specifies the replication state for a data type. If you specified multiple data types, the state applies to all the data types. For example, if you specified `DATATYPE=BACKUP, ARCHIVE`, the state applies to backup data and archive data.

The `REPLSTATE` parameter is optional. However, if you do not specify it, you must specify the `REPLRULE` parameter. You can specify one of the following values for the `REPLSTATE` parameter:

`Enabled`

Specifies that the data type is ready for replication.

`DISabled`

Specifies that replication does not occur until you enable it.

`PURGEdata`

Specifies that data is deleted from the target replication server. The type of data deleted is the type of data specified by the `DATATYPE` parameter. For example, if you specify `DATATYPE=BACKUP, ARCHIVE` and `REPLSTATE=PURGEDATA`, backup data and archive data are deleted from the file space on the target replication server.

After the data is deleted, the REPLSTATE parameter is set to DISABLED, preventing future replication of the data type or types. The replication rule for the data type is set to DEFAULT.

Remember: PURGEDATA processing does not delete file spaces. Only data is deleted. The file space shows as empty in the output of the QUERY OCCUPANCY command.

Example: Update replication rules for two data types

NODE1 has three file spaces: /a, /b, and /c. The replication rules for all file spaces are set to ALL_DATA. However, you want to replicate the backup and archive data in file space /a before the data in other file spaces is replicated.

```
update file space node1 /a datatype=backup,archive replrule=
    all_data_high_priority
```

Example: Update replication rules for two data types

NODE2 has two file spaces: /a and /b. You want to temporarily suspend replication of all data in file space /b.

```
update file space node2 /b datatype=backup,archive,spacemanaged
    replstate=disabled
```

Related commands

Table 1. Commands related to UPDATE FILESPACE

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET REPLRETENTION	Specifies the retention period for replication history records.
UPDATE NODE	Changes the attributes that are associated with a client node.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

UPDATE LIBRARY (Update a library)

Use this command to update a library definition.

AIX | **Windows** To update the device name, the ACS number, or the external manager path name of a library, you must use the UPDATE PATH command.

Linux To update the device name or the external manager path name of a library, you must use the UPDATE PATH command.

Syntax and parameter descriptions are available for the following library types.

- UPDATE LIBRARY (Update a 349X library)
- UPDATE LIBRARY (Update an ACSLS library)
- UPDATE LIBRARY (Update an EXTERNAL library)
- UPDATE LIBRARY (Update a FILE library)
- UPDATE LIBRARY (Update a manual library)
- UPDATE LIBRARY (Update a SCSI library)
- UPDATE LIBRARY (Update a shared library)
- UPDATE LIBRARY (Update a VTL library)

For detailed and current library support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Windows

To automatically label tape volumes in SCSI-type libraries, use the AUTOLABEL parameter on the DEFINE LIBRARY and UPDATE LIBRARY commands. Using this parameter eliminates the need to pre-label a set of tapes. It is also more efficient than using the LABEL LIBVOLUME command, which requires you to mount volumes separately. If you use the AUTOLABEL parameter, you must check in tapes by specifying CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

A label cannot include embedded blanks or periods and must be valid when used as a file name on the media.

You must label CD-ROM, Zip, or Jaz volumes with the device utilities from the manufacturer or the Windows utilities because IBM Spectrum Protect™ does not provide utilities to format or label these media types. The operating system utilities include the Disk Administrator program (a graphical user interface) and the label command.

Related commands

Table 1. Commands related to UPDATE LIBRARY

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE PATH	Changes the attributes associated with a path.

UPDATE LIBRARY (Update a 349X library)

Use this syntax to update a 349X library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRARY--library_name--+-----+----->
                                '-SHARed-----Yes---'
```

```

>----->
'-RESETDrives-----+Yes-+-'
                    '-No--'

>----->
'-AUTOLabel-----+No-----+'
                    +-Yes-----+
                    '-OVERWRITE-'

>-----<<
'-WORMSCRatchcategory----number-'

```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

SHAREd

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHAREd=YES to update a library that is not currently shared.

Important: If a library has a path from a data mover (such as a NAS file server) but no connection to the server, the library cannot be shared with another server.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

WORMSCRatchcategory

Specifies the category number to be used for WORM scratch volumes in the library. This parameter is required if you use WORM volumes. You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library. This parameter is only valid when 3592 WORM volumes are used.

Restriction: This parameter can only be updated if the device class WORM parameter is set to YES and the WORMSCRATCHCATEGORY currently has no defined value.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

The following table describes the three possible configurations for drives that are attached to NAS devices.

Table 1. Configurations for drives that are attached to NAS devices.

Library device configuration	The behavior for persistent reserve
------------------------------	-------------------------------------

Library device configuration	The behavior for persistent reserve
The library device is attached to the IBM Spectrum Protect server, and the tape drives are shared by the server and the NAS device.	Drive reservation preemption is supported when the NAS device supports persistent reserve and it is enabled. For more information about setting persistent reserve, see the documentation for your NAS device.
The library device is attached to the IBM Spectrum Protect server and the tape drives are accessed only from the NAS device.	Drive reservation preemption is not supported. If you enable persistent reserve on the NAS device for these drives and a reservation is set by the NAS device but never cleared, you must use another method to clear the reservation.

AIX | Windows

Yes

Specifies that drive preemption through persistent reserve or target reset are used.

No

Specifies that drive preemption through persistent reserve or target reset are not used. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used.

No

Specifies that drive preemption through persistent preserve is not used.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

Example: Add new devices to a shared library

Update a 3494 shared library named 3494LIB2 with new device names. AIX | Linux

```
update library 3494lib2 device=/dev/lmcp1,/dev/lmcp2,/dev/lmcp3
```

Windows

```
update library 3494lib device=lb3.0.0.0,lb4.0.0.0,lb5.0.0.0
```

UPDATE LIBRARY (Update an ACSLS library)

Use this syntax to update an ACSLS library.

Privilege class

Windows

In order to use ACSLS functions, the installation of StorageTek Library Attach software is required.

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>--UPDate LIBRary--library_name--+-----+----->
                                     '-SHARed-----Yes---'

>--+-----+----->
   '-RESEtDrives-----+-Yes+-'
                                     '-No--'

>--+-----+-----+-----+----->>
   '-AUTOLabel-----+-No-----+' '-ACSID-----number-'
                                     +-Yes-----+
                                     '-OVERWRITE-'
```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

SHARED

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

Important: If a library has a path from a data mover (such as a NAS file server) but no connection to the server, the library cannot be shared with another server.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

The following table describes the three possible configurations for drives that are attached to NAS devices.

Table 1. Configurations for drives that are attached to NAS devices.

Library device configuration	The behavior for persistent reserve
The library device is attached to the IBM Spectrum Protect server, and the tape drives are shared by the server and the NAS device.	Drive reservation preemption is supported when the NAS device supports persistent reserve and it is enabled. For more information about setting persistent reserve, see the documentation for your NAS device.
The library device is attached to the IBM Spectrum Protect server and the tape drives are accessed only from the NAS device.	Drive reservation preemption is not supported. If you enable persistent reserve on the NAS device for these drives and a reservation is set by the NAS device but never cleared, you must use another method to clear the reservation.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset are used.

No

Specifies that drive preemption through persistent reserve or target reset are not used. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used.

No

Specifies that drive preemption through persistent preserve is not used.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

ACSID (Required)

Specifies the number of this StorageTek library assigned by the ACSA (Automatic Cartridge System System Administrator). This can be a number from 0 to 126. Issue QUERY ACS on your system to get the number for your library ID. This parameter is required.

See your StorageTek documentation for more information.

Example: Update an ID number for an ACSLS library

Update an ACSLS library named ACSLSLIB with a new ID number.

```
update library acslslib acsid=1
```

UPDATE LIBRARY (Update an EXTERNAL library)

Use this syntax to update an external library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRARY--library_name----->
>--+-----+-----><
  '-AUTOLabel-----No-----'
                    +-Yes-----+
                    '-OVERWRITE-'
```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

Example: Update the path name for an external library

Update an external library named EXTLIB with a new path name for the media manager.

AIX | Linux

```
update library extlib externalmanager=/v/server/mediamanager
```

Windows

```
update library extlib externalmanager=c:\server\mediamanager
```

UPDATE LIBRARY (Update a FILE library)

Use this syntax to update a FILE library

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRARY--library_name--+-----+-----><
                                     '-SHARed-----Yes----'
```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

SHARed

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARed=YES to update a library that is not currently shared.

Important: If a library has a path from a data mover (such as a NAS file server) but no connection to the server, the library cannot be shared with another server.

Example: Update a FILE library to be shared

Update a file library named FILE2, so that it is shared:

```
update library file2 shared=yes
```

UPDATE LIBRARY (Update a manual library)

Use this syntax to update a manual library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRARY--library_name--+-----+----->
                                     '-RESEtDrives-----+Yes+-'
                                     '-No--'

>--+-----+-----><
  '-AUTOLabel-----+No-----+-'
                    +-Yes-----+
                    '-OVERWRITE-'
```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

RESEtDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect™ device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset are used.

No

Specifies that drive preemption through persistent reserve or target reset are not used. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used.

No

Specifies that drive preemption through persistent preserve is not used.

Note: A library manager is not able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

UPDATE LIBRARY (Update a SCSI library)

Use this syntax to update a SCSI library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRary--library_name----LIBType-----+-SCSI-+----->
                                     '-VTL--'
>--+-----+-----+-----+----->
  '-SHARed-----Yes---'  '-RESETDives-----+-Yes+-'
                                     '-No--'
>--+-----+-----+-----+----->
  '-AUTOLabel-----+-No-----+-'
                                     +-Yes-----+
                                     '-OVERWRITE-'
>--+-----+-----+-----+----->
  '-RELABELSCRatch-----+-No---+'
                                     '-Yes-'
>--+-----+-----+-----+----->>
  '-SERial-----+-serial_number+-'
```

Parameters

library_name (Required)

Specifies the name of the library to be updated.

LIBType (Required)

Specifies the library type that you want to update to. Possible values are:

VTL

Specifies that the library has a SCSI-controlled media changer device that is represented by a Virtual Tape Library. To mount volumes on drives in this type of library, IBM Spectrum Protect™ uses the media changer device. This value is effective when specified for libraries with a current library type of SCSI.

Note: Selecting the VTL library type assumes that the following conditions are true:

- Your environment does not include mixed-media
- Paths are defined between all drives in the library and all defined servers, including storage agents, that use the library

If both conditions are not met, performance can degrade to the same levels as the SCSI library type especially during times of high stress when most drives are in use concurrently.

SCSI

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, IBM Spectrum Protect uses the media changer device. This value is effective when specified for libraries with a current library type of VTL.

SHARED

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

Important: If a library has a path from a data mover (such as a NAS file server) but no connection to the server, the library cannot be shared with another server.

RESETDrives

Specifies whether the server preempts a drive reservation if the drive is already reserved by persistent reserve when the server tries to access the drive.

AIX | **Windows** If the drive is reserved by a SCSI-2 reserve, (and not by persistent reserve), the server uses a LUN reset to break the drive reservation to access the target device.

Linux LUN resets are not supported by the Linux operating system. If a drive is reserved by a SCSI-2 reserve, (and not by persistent reserve), the server is unable to break the reservation to access the drive. In this case, you can break the reservation by power cycling the device.

For Network-Attached Storage (NAS) devices, reservation is controlled by the NAS file server. IBM Spectrum Protect does not control NAS devices and the RESETDrives parameter is not relevant for NAS devices.

Support for persistent reserve has the following limitations:

- If you are using the IBM Spectrum Protect device driver, persistent reserve is supported only on some tape drives. For details, see Technote 1470319.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. For information about driver configuration, see the *IBM Tape Device Drivers Installation and User's Guide*.
- If you are using a virtual tape library that is emulating a supported drive, persistent reserve might not be supported.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset is used.

No

Specifies that drive preemption through persistent reserve or target reset is not used. The RESETDrives parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

- Yes
Specifies that drive preemption through persistent reserve is used.
- No
Specifies that drive preemption through persistent preserve is not used.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

- No
Specifies that the server does not attempt to label any volumes.
- Yes
Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

SERial

Specifies the serial number for the library being updated. This parameter is optional. The possible values are:

serial_number

Specifies the serial number for the library being updated.

If a path to this library has already been defined, then the number you enter here is compared to the number detected by IBM Spectrum Protect. If the numbers do not match, the command fails. If a path has not been defined, this serial number is verified when a path is defined.

AUTODetect

Specifies that the serial number is automatically detected and used by IBM Spectrum Protect if a path has already been defined to this library.

If a path to this library has not been defined, then the serial number is not detected.

RELABELSCRatch

Specifies whether the server relabels volumes that have been deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten. This parameter is optional and intended for use with a Virtual Tape Library (VTL).

Note: If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might affect performance.

- No
Specifies that the server does not relabel volumes that are deleted and returned to scratch.
- Yes
Specifies that the server relabels volumes that are deleted and returned to scratch.

UPDATE LIBRARY (Update a shared library)

Use this syntax to update a shared library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRary--library_name----->
>--PRIMarylibmanager---server_name-----<
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

PRIMarylibmanager

Specifies the name of the server that is responsible for controlling access to library resources. You must define this server with the DEFINE SERVER command before you can use it as a library manager.

Example: Change the library manager server for a library

For a library client server, change the name of the library manager server to CASTOR.

```
update library ltolib primarylibmanager=castor
```

UPDATE LIBRARY (Update a VTL library)

Use this syntax to update a library that is defined as VTL.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBRary--library_name----LIBType-----+VTL--+----->
                                     '-SCSI-'
>--+-----+-----+-----+-----+----->
  '-SHARed-----Yes---'  '-RESEtDrives-----+Yes--+-'
                                     '-No--'
>--+-----+-----+----->
  '-AUTOLabel-----+No-----+-'
                                     +-Yes-----+
                                     '-OVERWRITE-'
>--+-----+-----+----->
  '-RELABELSCRatch-----+No---+-'
                                     '-Yes-'
>--+-----+-----+----->>
  '-SERial-----+serial_number+--'
                                     '-AUTODetect----'
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType (Required)

Specifies the type of library that is being defined. Possible values are:

SCSI

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, IBM Spectrum Protect™ uses the media changer device. This value is effective when specified for libraries with a current library type of VTL.

VTL

Specifies that the library has a SCSI-controlled media changer device that is represented by a Virtual Tape Library. To mount volumes on drives in this type of library, IBM Spectrum Protect uses the media changer device. This value is effective when specified for libraries with a current library type of SCSI.

Note: Select the VTL library type only if the following conditions are true:

- Your environment does not include mixed-media
- Paths are defined between all drives in the library and all defined servers, including storage agents, that use the library

If both conditions are not met, performance can degrade to the same levels as the SCSI library type especially during times of high stress when most drives are in use concurrently.

SHARED

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

Important: If a library has a path from a data mover (such as a NAS file server) but no connection to the server, the library cannot be shared with another server.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

AIX | **Windows** If persistent reserve is not supported, the server completes a reset of the path to the target device.

Linux If persistent reserve is not supported, the server is not able to reset the path to the target device.

Support for persistent reservation has the following limitations:

- If you are using the IBM Spectrum Protect device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 for details.
- If you are using the IBM® device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

AIX | **Windows**

Yes

Specifies that drive preemption through persistent reserve or target reset are used.

No

Specifies that drive preemption through persistent reserve or target reset are not used. The RESETDRIVES parameter must be set to YES in a clustered environment when SHARED=NO.

Linux

Yes

Specifies that drive preemption through persistent reserve is used.

No

Specifies that drive preemption through persistent preserve is not used.

Note: A library manager is not able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the CHECKIN LIBVOLUME command.

No

Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RELABELSCRatch

Specifies whether the server relabels volumes that have been deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten.

Note: If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might affect performance.

Yes

Specifies that the server relabels volumes that are deleted and returned to scratch.

No

Specifies that the server does not relabel volumes that are deleted and returned to scratch.

SERial

Specifies the serial number for the library being updated. This parameter is optional. The possible values are:

serial_number

Specifies the serial number for the library being updated.

If a path to this library has already been defined, then the number you enter here is compared to the number detected by IBM Spectrum Protect. If the numbers do not match, then the command fails. If a path has not been defined, this serial number is verified when a path is defined.

AUTODetect

Specifies that the serial number is automatically detected and used by IBM Spectrum Protect if a path has already been defined to this library.

If a path to this library has not been defined, then the serial number is not detected.

UPDATE LIBVOLUME (Change the status of a storage volume)

Use this command to change the status of a sequential access storage volume in a library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate LIBVolume--library_name--volume_name--STATus-----+PRIVate+--->
                                     '-SCRatch-'
>--+-----+-----><
   '-OWNer-----server_name-'
```

Parameters

library_name (Required)

Specifies the name of the library.

volume_name (Required)

Specifies the volume name of the storage volume.

STATus (Required)

Specifies a change to the status of a storage volume. Possible values are as follows:

PRIVate

Specifies that the server updates the storage volume to a private volume.

SCRatch

Specifies that the server updates the storage volume to a scratch volume.

Restriction: You cannot change the status of a volume from private to scratch if the volume belongs to a storage pool or is defined in the volume history file. You can change the status if you make a mistake when you check in volumes to the library and assign the volumes the wrong status.

AIX	Linux	Windows	OWNer
-----	-------	---------	-------

Specifies which server owns a private volume in a shared library that is shared across a SAN. You can change the owner of a private volume in a shared library (SAN) when you issue the command from the library manager server. If you do not specify this parameter, the library manager server owns the private volume.

Important: Do not use OWNER as a value for scratch volumes. However, you can use OWNER when you change a scratch volume to private.

Example: Update a volume's status

Update the volume that is named WPDV00 in the library that is named AUTO to reflect a status of PRIVATE.

```
update libvolume auto wpdv00 status=private
```

Related commands

Table 1. Commands related to UPDATE LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
AIX Linux Windows LABEL LIBVOLUME	AIX Linux Windows Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.

UPDATE MACHINE (Update machine information)

Use this command to update machine information. This information will be included in the plan file to help you to recover the client machines.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate MACHINE--machine_name----->
>--+-----+-----+-----+----->
  '-DESCRiption----description-'  '-BUilding----building-'
>--+-----+-----+-----+----->
  '-FLoor----floor-'  '-ROom----room-'
>--+-----+-----+-----+-----><
  '-PRIority----number-'  '-ADSMServer-----+Yes+-'
                               '-No--'
```

Parameters

machine_name (Required)

Specifies the name of the machine to be updated.

DESCRiption

Specifies a description of the machine. This parameter is optional. The text can be up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

BUilding

Specifies the name or number of the building that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

FLoor

Specifies the name or number of the floor that this machine is on. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

ROom

Specifies the name or number of the room that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

PRIority

Specifies the restore priority for the machine as an integer from 1 to 99. The highest priority is 1. This parameter is optional. Use this value to prioritize client machine recovery.

ADSMServer

Specifies whether the machine contains an IBM Spectrum Protect™ server. This parameter is optional. Possible values are:

No

This machine does not contain an IBM Spectrum Protect server.

Yes

This machine contains an IBM Spectrum Protect server. Only one machine can be defined as containing an IBM Spectrum Protect server.

Example: Update information for a specific machine

Update the DISTRICT5 machine information to reflect that it contains the server.

```
update machine district5 admsserver=yes
```

Related commands

Table 1. Commands related to UPDATE MACHINE

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DELETE MACHINE	Deletes a machine.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the IBM Spectrum Protect database.
QUERY MACHINE	Displays information about machines.

UPDATE MGMTCLASS (Update a management class)

Use this command to change a management class. To allow clients to use the updated management class, you must activate the policy set that contains the management class.

Important: The UPDATE MGMTCLASS command fails if a copy storage pool is specified as the destination for files that were migrated by an IBM Spectrum Protect™ for Space Management client.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-UPDate MGmtclass--domain_name--policy_set_name--class_name--->
>--+-----+----->
  '-SPACEMGTEchnique-----+AUTOMATIC++'
                                +-SElective+
                                '-NONE-----'
>--+-----+----->
  '-AUTOMIGNonuse-----days-'
>--+-----+----->
  '-MIGREQUIRESBkup-----+Yes++'
                                '-No--'
>--+-----+----->
  '-MIGDESTination-----pool_name-'
>--+-----+-----><
```

'-DESCRiption-----description-'

Parameters

domain_name (Required)

Specifies the policy domain to which the management class belongs.

policy_set_name (Required)

Specifies the policy set to which the management class belongs. You cannot update a management class that belongs to the ACTIVE policy set.

class_name (Required)

Specifies the management class to update.

SPACEMGTECHnique

Specifies whether a file using this management class is eligible for migration. This parameter is optional. This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients. Possible values are:

AUTOMATIC

Specifies that the file is eligible for both automatic migration and selective migration.

SELECTive

Specifies that the file is eligible for selective migration only.

NONE

Specifies that the file is not eligible for migration.

AUTOMIGNonuse

Specifies the number of days that must elapse since a file was last used before it is eligible for automatic migration. This parameter is optional. If SPACEMGTECHNIQUE is not AUTOMATIC, the server ignores this attribute. You can specify an integer from 0 to 9999.

This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients.

MIGREQUIRESBkup

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients. Possible values are:

Yes

Specifies that a backup version must exist.

No

Specifies that a backup version is optional.

MIGDESTination

Specifies the primary storage pool where the server initially stores files migrated by IBM Spectrum Protect for Space Management clients. This parameter is effective only for IBM Spectrum Protect for Space Management clients, not for backup-archive clients or application clients.

The command fails if you specify a copy storage pool as the destination.

DESCRiption

Specifies a description of the management class. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("").

Example: Update the policy domain and storage pool of a specific management class

For the management class ACTIVEFILES, in policy set VACATION in the EMPLOYEE_RECORDS policy domain, change the storage pool where migrated files are stored.

```
update mgmtclass employee_records vacation
activefiles migdestination=diskpool2
```

Related commands

Table 1. Commands related to UPDATE MGMTCLASS

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

UPDATE NODE (Update node attributes)

Use this command to modify the attributes of a registered node.

You must use the RENAME NODE command to change the name of a registered node.

If you update the node authentication method or the node SSLREQUIRED setting and there is a same-named administrator, those administrator ID settings change.

You must have system level authority to update the node authentication method or the node SSLREQUIRED setting and also update a same-named administrator ID. If the same-named administrator ID has client owner authority over the node that is being updated, then system level authority is not required. You must have either unrestricted policy privilege or restricted policy privilege for the policy domain to which the client node belongs.

For users of Lightweight Directory Access Protocol (LDAP) servers:

- The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see Managing passwords and logon procedures.
- If you change the authentication mode to LDAP, and the node name matches an administrative user ID, you might see unexpected behavior when an automatic password change occurs because the password might be updated twice. As a result, the password might become unknown to the administrative user ID. Alternatively, the password update operation might fail.

When you register or update a node, you can specify whether damaged files on the node can be recovered from a target replication server. Files can be recovered only if all the following conditions are met:

- Version 7.1.1 or later, is installed on the source and target replication servers.
- The REPLRECOVERDAMAGED system parameter is set to ON. The system parameter can be set by using the SET REPLRECOVERDAMAGED command.
- The source server includes at least one file that is marked as damaged in the node that is being replicated.
- The node data was replicated before the damage occurred.

The following table describes how parameter settings affect the recovery of damaged, replicated files.

Table 1. Settings that affect the recovery of damaged files

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result

Setting for the REPLRECOVERDAMAGED system parameter	Value of the RECOVERDAMAGED parameter on the REPLICATE NODE command	Value of the RECOVERDAMAGED parameter on the REGISTER NODE and UPDATE NODE commands	Result
OFF	YES, NO, or not specified	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
OFF	ONLY	YES or NO	An error message is displayed because files cannot be recovered when the REPLRECOVERDAMAGED system parameter is set to OFF.
ON	YES	YES or NO	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	NO	YES or NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.
ON	ONLY	YES or NO	Damaged files are recovered from the target replication server, but standard node replication does not occur.
ON	Not specified	YES	During node replication, standard replication occurs and damaged files are recovered from the target replication server.
ON	Not specified	NO	During node replication, standard replication occurs and damaged files are not recovered from the target replication server.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax

```

(1)
>>-UPDate Node-----node_name----->
>+-----+-----+-----+-----+----->
| (2) |
+-----+-----+-----+-----+-----+
|      password--+-----+-----+-----+-----+
|      '-FORCEPwreset-----+No--+-' |
|      '-Yes-' |
| '-FORCEPwreset-----Yes-----' |
>+-----+-----+-----+-----+----->
| '-PASSExp-----days-' | '-CLOptset-----option_set_name-'
>+-----+-----+-----+-----+----->
| '-CONtact-----text-' | '-DOMain-----domain_name-'
>+-----+-----+-----+-----+----->
| '-COMPrESSION-----+Client--+-' | '-ARCHDElete-----+Yes--+-'
|      +-Yes-----+ |      '-No--'
|      '-No-----'

```

```

>----->
'-BACKDElete-----+No--+'
      '-Yes-'

>----->
'-WHEREDomain-----domain_name-'

>----->
'-WHEREPlatform-----client_platform_name-'

>----->
'-MAXNUMMP-----number-' '-KEEPMP-----+No--+'
      '-Yes-'

>----->
'-URL-----url_address-' '-UTILITYUrl-----utility_url-'

(3)
>----->
'-AUTOFSRename-----+Yes-----+'
      +-No-----+
      '-Client-'

>----->
'-VALIDateprotocol-----+No-----+'
      +-Dataonly+
      '-All-----'

>----->
'-TXNGroupmax-----+0-----+'
      '-number-'

.-DATAWritepath-----ANY-----
>----->
'-DATAWritepath-----+ANY-----+'
      +-LAN-----+
      '-LANFree-'

.-DATAReadpath-----ANY-----
>----->
'-DATAReadpath-----+ANY-----+'
      +-LAN-----+
      '-LANFree-'

>----->
'-TARGETLevel-----V.R.M.F-'

.-SESSIONINITiation-----Clientorserver-----
>----->
'-SESSIONINITiation-----+Clientorserver-----+'
      |
      '-SERVEROnly--HLAddress-----ip_address--LLAddress-----tcp_port-----'
      (4) |

>----->
'-HLAddress-----ip_address-'

>----->
|
| (4) |
'-LLAddress-----tcp_port-----'

>----->
'-EMAILAddress-----userID@node-'

>----->
'-DEDUPlication-----+SERVEROnly-----+'
      '-Clientorserver-'

>----->
|
| (5) |
'-BACKUPINITiation-----+All-----+'
      '-ROOT-'

>----->

```


8. The SYNCLDAPDELETE parameter applies only if a node that authenticates to a Lightweight Directory Access Protocol (LDAP) server reverts to local authentication.
9. The SSLREQUIRED parameter is deprecated.

Parameters

node_name (Required)

Specifies the name of the client node to be updated. You can use wildcard characters to specify this name.

Restriction: When you update a password with the UPDATE NODE command, you cannot use a wildcard character with the node_name parameter.

password

Specifies the new password for the client node. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters. This parameter is optional in most cases. If the node authentication method is changed from LDAP to LOCAL, a password is required. If the node authentication method is LDAP, do not specify a password by using the UPDATE NODE command. Passwords remain current for a period that is determined by the password expiration period.

FORCEPwreset

Specifies whether to force a client to change or reset the password. This parameter is optional. You can specify one of the following values:

No

Specifies that the password expiration period is set by the SET PASSEXP command. Do not force a client to change or reset the password while it attempts to log on to the server.

Yes

Specifies that the client node or administrator password will expire at the next logon. The client must change or reset the password at the next logon.

Restrictions:

- For nodes that authenticate with an LDAP server, password expiration is set by using LDAP server utilities. For this reason, do not specify FORCEPWRESET=YES if you plan to specify AUTHENTICATION=LDAP.
- If you plan to update a node to authenticate with an LDAP server, and you specified FORCEPWRESET=YES, you must change the password before you can specify FORCEPWRESET=NO and AUTHENTICATION=LDAP.

PASSExp

Specifies the number of days the password remains valid. You can set the password expiration period in the range 0 - 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the password expiration period is unchanged.

You can change the password expiration period by using the UPDATE NODE or SET PASSEXP commands. To set a common expiration period for all administrators and client nodes, issue the SET PASSEXP command. You can also use the SET PASSEXP command to selectively set password expiration periods. If you selectively set a password expiration period by using the REGISTER NODE command, the UPDATE NODE command, or the SET PASSEXP command, the expiration period is excluded from common password expiration periods that were created by using the SET PASSEXP command.

You can use the RESET PASSEXP command to reset the password expiration period to the common expiration period. This parameter does not apply to passwords that authenticate with an LDAP directory server.

CLOptset

Specifies the name of the option set to be used by the client. This parameter is optional. To remove a client option set, specify the CLOPTSET parameter with a null string ("").

CONtact

Specifies a text string of information that identifies the client node. This parameter is optional. The maximum length of the text string is 255 characters. Enclose the contact information in quotation marks if it contains any blanks. To remove previously defined contact information, specify a null string ("").

DOmain

Specifies the name of the policy domain to which you want to register the client node. This parameter is optional.

Restriction: For servers with data retention protection enabled, an archived registered node cannot be reassigned to a different policy domain.

COMPression

Specifies whether the client node compresses its files before it sends them to the server for backup and archive. This parameter is optional.

Restriction: This parameter cannot be specified for a NAS node.

You can specify one of the following values:

Client

Specifies that the client determines whether files are to be compressed.

Yes

Specifies that the client node compresses its files before it sends them to the server for backup and archive.

No

Specifies that the client node does not compress its files before it sends them to the server for backup and archive.

ARCHDElete

Specifies whether the client node can delete its own archived files from the server. This parameter is optional. You can specify one of the following values:

Yes

Specifies that the client node can delete its own archive files from the server.

No

Specifies that the client node cannot delete its own archive files from the server.

BACKDElete

Specifies whether the client node can delete its own backup files from the server. This parameter is optional. You can specify one of the following values:

No

Specifies that the client node cannot delete its own backup files from the server.

Yes

Specifies that the client node can delete its own backup files from the server.

WHEREDomain

Specifies the name of the policy domain to be used as a filter in combination with the node name to select nodes to update. This parameter is optional.

WHEREPlatform

Specifies the name of the client platform to be used as a filter in combination with the node name to select nodes to update. This parameter is optional.

MAXNUMMP

Specifies the maximum number of mount points a node can use on the server or storage agent only for operations such as backup, archive, and IBM Spectrum Protect for Space Management migration. The parameter is optional and does not apply to nodes with a type of NAS or SERVER. The default value is 1. You can specify an integer in the range 0 - 999. A value of 0 specifies that a node cannot acquire any mount point for a client data store operation. The MAXNUMMP value is not evaluated or enforced during client data read operations such as restore, retrieve, and IBM Spectrum Protect for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node and might prevent the data store operations from being able to acquire mount points.

For volumes in a storage pool that is associated with the FILE or CENTERA device type, the server can have multiple sessions to read and one process to write to the same volume concurrently. To increase concurrency and provide efficient access for nodes with data in FILE or CENTERA storage pools, increase the value of the MAXNUMMP parameter.

For nodes that store data into primary storage pools with the simultaneous-write function that is enabled, you must adjust the value of the MAXNUMMP parameter to specify the correct number of mount points for each client session. A client session requires one mount point for the primary storage pool and one mount point for each copy storage pool and each active-data pool.

URL

Specifies the URL of the IBM Spectrum Protect web client that is configured on the client system. You can use the URL in a web browser and in the Operations Center to remotely manage the client node.

This parameter is optional. The URL must include the DNS name or IP address of the client system, and the port number that is defined on the client system for the IBM Spectrum Protect web client. For example,
`http://client.mycorp.com:1581`

If you want to remove the value from this parameter, specify empty single quotation marks or empty double quotation marks with no spaces (" for single quotation marks, or "" for double quotation marks).

UTILITYUrl

Specifies the address of the IBM Spectrum Protect client management services that are configured on the client system. This URL is used by the Operations Center to access client log files so that you can remotely diagnose client issues from the Operations Center.

This parameter is optional. You can specify a URL of up to 200 characters in length. The URL must start with `https`. It includes the DNS name or IP address of the client system, and the port number that is defined on the client system for the IBM Spectrum Protect client management services. For example, `https://client.mycorp.com:9028`

If you omit the port number, the Operations Center uses the port number 9028, which is the default port number when you install the client management services on the client system.

KEEPMP

Specifies whether the client node keeps the mount point for the entire session. The parameter is optional. You can specify one of the following values:

No

Specifies that the client node releases the mount point during the session. If policy definitions cause data to be stored to a disk storage pool after data is stored to a sequential access storage pool, any mount points that are held by the session will be released.

Yes

Specifies that the client node must retain the mount point during the entire session. If policy definitions cause data to be stored to a disk storage pool after data is stored to a sequential access storage pool, any mount points that are held by the session will not be released.

AUTOFSRename

Specifies whether the client is prompted for renaming file spaces when the client system upgrades to a client that supports Unicode. The prompting and renaming, if allowed, occur only when the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming changes the names of existing backed-up file spaces that are not in Unicode in server storage. Then, the file spaces are backed up in Unicode. You can use this parameter for Unicode-enabled IBM Spectrum Protect clients by using Windows, Macintosh OS X, and NetWare operating systems.

Important: After the client with support for Unicode is installed, any new file spaces that the client backs up are stored in server storage by using the UTF-8 code page. UTF-8 is a byte-oriented encoding form that is specified by the Unicode Standard.

You can specify one of the following values:

Yes

The server automatically renames existing file spaces when the client system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming occurs whether the client uses the graphical user interface, the command line, or the client scheduler.

For example, the server renames a drive as follows:

- Original name: D_DRIVE
- New name: D_DRIVE_OLD

The new name indicates that the file space is stored on the server in format that is not Unicode.

No

The server does not rename file spaces automatically when the client system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup.

Client

The option `AUTOFSRENAME` in the client option file determines whether file spaces are renamed.

By default, the client option is set to `PROMPT`. When the client system upgrades to a client that supports Unicode and the client runs an IBM Spectrum Protect operation with the graphical user interface or the command line, the program displays a one-time prompt to the user about whether to rename file spaces.

When the client scheduler runs an operation, the program does not prompt for a choice about renaming, and does not rename file spaces. Backups of existing file spaces are sent as before (not in Unicode).

VALIDateprotocol (deprecated)

Specifies whether IBM Spectrum Protect performs a cyclic redundancy check to validate the data that is sent between the client and the server. The parameter is optional.

Important: Beginning with IBM Spectrum Protect Version 8.1.2 and Tivoli® Storage Manager Version 7.1.8, this parameter is deprecated. The validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The VALIDATEPROTOCOL parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

TXNGroupmax

Specifies the number of files that are transferred as a group between a client and a server between transaction commit points. Client performance might be improved by using a larger value for this option.

Specifying 0 indicates that the node uses the server global value that is set in the server options file. To use a value other than the server global value, specify a value of 4 through 65,000 for this parameter. The node value takes precedence over the server value.

Tip: Increasing the TXNGROUPMAX value increases recovery log utilization. Higher recovery log utilization might increase the risk of running out of log space. Evaluate the performance of each node before you change the parameter.

DATAWritepath

Specifies the transfer path that is used when the client sends data to the server, storage agent, or both, during storage operations such as backup or archive. The parameter is optional.

Remember: If a path is unavailable, the node cannot send any data. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation fails.

You can specify one of the following values:

ANY

Specifies that data is sent to the server, storage agent, or both, using any available path. A LAN-free path is used if one is available. If a LAN-free path is unavailable, the data is moved over the LAN.

LAN

Specifies that data is sent over the LAN.

LANFree

Specifies that data is sent over a LAN-free path.

DATAReadpath

Specifies the transfer path that is used when the server, storage agent, or both read data for a client, during operations such as restore or retrieve. The parameter is optional.

Remember: If a path is unavailable, data cannot be read. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation fails. The value for the transfer path also applies to failover connections. If the value is set to LANFree, failover cannot occur for the node on the secondary server.

You can specify one of the following values:

ANY

Specifies that the server, storage agent, or both use any available path to read data. A LAN-free path is used if one is available. If a LAN-free path is unavailable, the data is read over the LAN.

LAN

Specifies that data is read over the LAN.

LANFree

Specifies that data is read by using a LAN-free path.

SESSIONINITiation

Controls whether the server or the client initiates sessions. The parameter is optional.

Clientorserver

Specifies that the client might initiate sessions with the server by communicating on the TCP/IP port that is defined with the server option TCPPOINT. Server-prompted scheduling might also be used to prompt the client to connect to the server.

SERVEROnly

Specifies that the server does not accept client requests for sessions. All sessions must be initiated by server-prompted scheduling on the port that is defined for the client with the REGISTER or UPDATE NODE commands. You cannot use the client acceptor, dsmcad, to start the scheduler when SESSIONINITIATION is set to SERVERONLY.

HLAddress

Specifies the client IP address that the server contacts to initiate scheduled events. This parameter must be used when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that are previously used by the client to contact the server.

The address can be specified either in numeric or host name format. If a numeric address is used, it is saved without verification by a domain name server. If the address is not correct, it can cause failures when the server attempts to contact the client. Host name format addresses are verified with a domain name server. Verified names are saved and resolved with Domain Name Services when the server contacts the client.

LLAddress

Specifies the client port number on which the client listens for sessions from the server. This parameter must be used when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that are previously used by the client to contact the server.

The value for this parameter must match the value of client option TCPCLIENTPORT. The default value is 1501.

HLAddress

Specifies the client IP address that the server contacts to initiate scheduled events. This optional parameter is used only when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that were previously used by the client to contact the server. If SESSIONINITIATION SERVERONLY is not in use, this option has no effect.

The address can be specified either in numeric or host name format. If a numeric address is used, it is saved without verification by a domain name server. If the address is not correct, it can cause failures when the server attempts to contact the client. Host name format addresses are verified with a domain name server. Verified names are saved and resolved with Domain Name Services when the server contacts the client.

LLAddress

Specifies the client port number on which the client listens for sessions from the server. This optional parameter is used only when SESSIONINITIATION is set to SERVERONLY, regardless of any addresses that were previously used by the client to contact the server. If SESSIONINITIATION SERVERONLY is not in use, this option has no effect.

The value for this parameter must match the value of client option TCPCLIENTPORT. The default value is 1501.

EMAILAddress

This parameter is used for more contact information. The information that is specified by this parameter is not acted upon by IBM Spectrum Protect.

DEDUPLICATION

Specifies where data deduplication can occur for this node. You can specify one of the following values:

SERVEROnly

Specifies that data that is stored by this node can be deduplicated on the server only.

Clientorserver

Specifies that data that is stored by this node can be deduplicated on either the client or the server. For data deduplication to take place on the client, you must also specify a value of YES for the DEDUPLICATION client option. You can specify this option in the client option file or in the client option set on the IBM Spectrum Protect server.

TARGETLevel

Specifies the client deployment package that is targeted for this node. You can substitute an applicable release package for V.R.M.F (Version.Release.Modification.Fix) Level. For example: TARGETLevel=6.2.0.0.

You must specify each segment with a number that is applicable to a deployment package. You cannot use an asterisk in any field as a substitution for a valid number. To remove an existing value, specify a null string (" "). The parameter is optional.

Restriction: The TARGETLEVEL parameter does not apply to nodes with a type of NAS or SERVER.

BACKUPINITiation

Specifies whether the non-root user ID on the client node can back up files to the server. The parameter is optional. The default value is ALL, indicating that non-root user IDs can back up data to the server. You can select one of the following values:

All

Specifies that non-root user IDs can back up files to the server. ALL is the default if BACKUPINITIATION is not specified.

ROOT

Specifies that only the root user ID can back up files to the server.

Restriction: The attribute is ignored by the server if the backup-archive client connects from an operating system other than AIX®, Linux, or Mac OS.

BKREPLRuledefault, ARREPLRuledefault, and SPREPLRuledefault

Specifies the replication rule that applies to a data type if the file space rules for the data type are set to DEFAULT:

BKREPLRuledefault

Specifies the replication rule for backup data.

ARREPLRuledefault

Specifies the replication rule for archive data.

SPREPLRuledefault

Specifies the replication rule for space-managed data.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that a client node contains active backup data and archive data. Replication of the active backup data is a higher priority than the archive data. To prioritize both types of data, specify

`BKREPLRULEDEFAULT=ACTIVE_DATA_HIGH_PRIORITY ARREPLRULEDEFAULT=ALL_DATA`.

You can specify the following rules:

ALL_DATA

Replicates active and inactive backup data, archive data, or space-managed data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority. This rule is valid only for BKREPLRULEDEFAULT.

Attention:

If you specify ACTIVE_DATA and one or more of the following conditions are true, inactive backup data on the target replication server is deleted, and inactive backup data on the source replication server is not replicated.

- When a release version earlier than Version 7.1.1 is installed on either the source or target replication servers.
- When you are using the REPLICATE NODE command with the `FORCERECONCILE=YES` parameter.
- When you are running the initial replication of a file space after you configure replication, restore the database, or upgrade both the source and target replication servers from a release version earlier than V7.1.1.

If the previous conditions are not true, all new and changed files since the last replication are replicated, including inactive files, and files are deleted when they expire.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data, archive data, or space-managed data. Data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

This rule is the same as the ACTIVE_DATA replication rule except data is replicated with a high priority. This rule is valid only for BKREPLRULEDEFAULT.

DEFAULT

Replicates data according to the server replication rule for backup data.

For example, suppose that you want to replicate the archive data in all the file spaces that belongs to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify

`ARREPLRULEDEFAULT=DEFAULT`. Ensure that the file space rules for archive data are also set to DEFAULT and that the server rule for archive data is set to ALL_DATA_HIGH_PRIORITY.

Restriction: If a node is configured for replication, the file space rules are set to DEFAULT after the node stores data on the source replication server.

NONE

Data of the specified type is not replicated.

For example, if you do not want to replicate space-managed data that belongs to a client node, specify
`SPREPLRULEDEFAULT=NONE`

REPLState

Specifies whether data that belongs to the client node is ready to be replicated. This parameter is optional. You can specify one of the following values:

ENabled

Specifies that the client node is ready for replication.

DISabled

Specifies that replication does not occur until you enable it.

The system response to these settings depends on the following factors:

Whether the client node definition exists only on the source replication server and you are configuring the client node for replication for the first time

If you set the replication state to ENABLED or DISABLED, the replication mode of the node on the source replication server is automatically set to SEND after the UPDATE NODE command is issued. When replication first occurs, a client node definition on the target server is automatically created. The replication state of the client node on the target server is automatically set to ENABLED. The replication mode is set to RECEIVE.

Whether the client node definition exists on the source and the target replication servers, and the node data was previously replicated

For replication to occur, the replication state of the client node on both the source and the target servers must be set to ENABLED. For example, if the replication state of a client node on the source server is ENABLED and the replication state on the target server is DISABLED, replication does not occur.

Whether the client node definition exists on the source and the target replication servers, and the node data was previously exported from the source replication server and imported to the target replication server

In this case, you are configuring the client nodes to synchronize the data between the two servers. When replication first occurs, the replication state of the client node on the target server is automatically set to ENABLED. Data on the source and target servers is synchronized.

Restriction: To synchronize data, you must specify the REPLMODE parameter in addition to the REPLSTATE parameter.

You can specify the REPLMODE parameter only if the client node has never been replicated:

- If the client node definition exists only on the source replication server, the replication mode of the node on the source replication server is automatically set to SEND when the UPDATE NODE command is issued. The replication mode of the node on the target replication server is automatically set to RECEIVE.
- If data that belongs to the node was previously replicated, the replication mode of the node on the source replication server is SEND. The replication mode of the node on the target replication server is RECEIVE.

REPLMode

Specifies whether to synchronize the data that belongs to this client node. Specify this parameter only if data that belongs to the client node was exported from the source replication server and imported to the target replication server. Synchronization occurs during replication.

To synchronize data, you must issue the UPDATE NODE command on both the source and target replication servers and specify the REPLMODE and REPLSTATE parameters. The value that you specify for the REPLMODE parameter depends on whether the server is a source of or a target for replicated data.

You can specify one of the following values:

SYNCSEnd

Specifies that data that belongs to this client node is synchronized with data on a target server during replication. Specify this value only on the server that exported the data. When the synchronization is complete, the replication mode for the client node on the source server is automatically set to SEND. The replication mode remains SEND unless you remove the node by issuing the REMOVE REPLNODE command.

SYNCRECeive

Specifies that data that belongs to this client node is synchronized with data on a source server during replication. Specify this value only on the server that imported the data. When the synchronization is complete, the replication

mode for the client node on the target server is automatically set to RECEIVE. The replication mode remains RECEIVE unless you remove the node by issuing the REMOVE REPLNODE command.

Restrictions:

- You can set the REPLMODE parameter only if the initial replication state is NONE. To synchronize data, you change the replication state to ENABLED or DISABLED and specify a value for the REPLMODE parameter.
- Data can be synchronized only if you specified DATES=ABSOLUTE on the IMPORT NODE command. If you specified DATES=RELATIVE to import data, you must rename the node or delete its data before replication. If you do not take one of these steps, you can lose data.
- If the REPLMODE parameter was set incorrectly, you must issue the REMOVE REPLNODE command before you update the client node definition. For example, suppose that you updated the definition of a client node whose data you wanted to replicate. The data that belongs to the node was previously exported to the target replication server. You specified ENABLED as the setting of the REPLSTATE parameter. However, you did not specify SYNCSEND on the source replication server. As a result, the REPLMODE parameter was automatically set to SEND, and data that belongs to the node could not be synchronized or replicated.

Issuing REMOVE REPLNODE sets the replication state and the replication mode to NONE. After the REMOVE REPLNODE command is completed, reissue the UPDATE NODE command with the correct parameters and values.

RECOVERDamaged

Specifies whether damaged files can be recovered for this node from a target replication server. The parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that recovery of damaged files from a target replication server is enabled for this node.

No

Specifies that recovery of damaged files from a target replication server is not enabled for this node.

Tip: The value of the RECOVERDAMAGED parameter is only one of several settings that determine whether damaged files are recovered. For information about how to specify the settings, see Settings that affect the recovery of damaged files.

ROLEOVERRIDE

Specifies whether to override the reported role of the client for processor value unit (PVU) estimation reporting. The default is USERREPORTED.

The role reported by the client is either client-device (for example, a workstation) or server-device (for example, file/print server, application server, database). By default, the client reports its role that is based on the client type and the operating system. All clients initially report their role as server-device, except for IBM Spectrum Protect backup-archive clients that are running Microsoft Windows workstation distributions (Windows Vista) and Macintosh OS X.

Specify one of the following values:

Client

Specifies a client-device.

Server

Specifies a server-device.

Other

Specifies that this node is not to be used for PVU estimation reporting. The Other value is useful when multiple nodes are deployed for a physical system (for example, virtual environments, test nodes, retired nodes, and nodes not in production or clustering).

Usereported

Use the reported role that is provided by the client.

AUTHentication

This parameter determines the password authentication method that you use; either LDAP or LOCAL.

Local

Specifies that the node uses the local IBM Spectrum Protect server database to store passwords.

LDap

Specifies that the node uses an LDAP directory server to authenticate passwords. Passwords are not stored in the IBM Spectrum Protect database.

SYNCLdapdelete

This parameter applies only if you want a node that authenticates with a Lightweight Directory Access Protocol (LDAP) server to change to authenticate with the IBM Spectrum Protect server. The parameter specifies whether to remove the node from the LDAP server.

Yes

Specifies that the node is removed.

Restriction: Do not specify a value of YES. (The value of YES is appropriate only for users of the previous LDAP authentication method, which is described in Managing passwords and logon procedures.)

No

Specifies that the node is not removed. This is the default value.

SSLrequired (deprecated)

Specifies whether the node must use the Secure Sockets Layer (SSL) protocol to communicate with the IBM Spectrum Protect server. The parameter is optional. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Important: Beginning with IBM Spectrum Protect V8.1.2 software and Tivoli Storage Manager V7.1.8 software, this parameter is deprecated. The validation that was enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The SSLREQUIRED parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

SESSIONSECURITY

Specifies whether the node must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRICT

Specifies that the strictest security settings are enforced for the node. The STRICT value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the server and the node. To specify whether the server uses TLS 1.2 for the entire session or only for authentication, see the SSL client option.

To use the STRICT value, the following requirements must be met to ensure that the node can authenticate with the server:

- Both the node and server must be using IBM Spectrum Protect software that supports the SESSIONSECURITY parameter.
- The node must be configured to use the TLS 1.2 protocol for SSL sessions between the server and the node.

Nodes set to STRICT that do not meet these requirements are unable to authenticate with the server.

TRANSITIONAL

Specifies that the existing security settings are enforced for the node. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the STRICT value.

If SESSIONSECURITY=TRANSITIONAL and the node has never met the requirements for the STRICT value, the node will continue to authenticate by using the TRANSITIONAL value. However, after a node meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. Then, the node can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT. In addition, after a node successfully authenticates by using a more secure communication protocol, the node can no longer authenticate by using a less secure protocol. For example, if a node that is not using SSL is updated and successfully authenticates by using TLS 1.2, the node can no longer authenticate by using no SSL protocol or by using TLS 1.1. This restriction also applies when you use functions such as virtual volumes, when the node authenticates to the IBM Spectrum Protect server as a node from another server.

SPLITLARGEObjects

Specifies whether large objects that are stored by this node are automatically split into smaller pieces, by the server, to optimize server processing. Specifying Yes causes the server to split large objects (over 10 GB) into smaller pieces when stored by a client node. Specifying No bypasses this process. Specify No only if your primary concern is maximizing throughput of backups directly to tape. The default value is Yes.

Example: Update node SIMON to authenticate with an LDAP directory server and connect using SSL

```
update node simon authentication=ldap sslrequired=yes
```

When you specify the SSLREQUIRED parameter, the server is not automatically configured for SSL. You must follow the instructions for connecting with SSL in order for the example to work.

Example: Update all nodes to communicate with a server by using strict session security

Update all nodes to use the strictest security settings to authenticate with the server.

```
update node * sessionsecurity=strict
```

Example: Update a node with software release information for a future deployment

The client deployment feature helps you update a backup-archive client to a newer release. The information that is generated from the UPDATE NODE command can help you when you plan a deployment. The information is stored for a future deployment and can be viewed by issuing the QUERY NODE command. After a deployment, you can issue the QUERY NODE command to see the current level and the target level. For example, to update node LARRY to backup-archive client Version 6.3.0.0.

```
update node LARRY targetlevel=6.3.0.0
```

Example: Update a node backup to compress data and keep the client from deleting archived files

Update node LARRY so that the data on node LARRY is compressed when it is backed up or archived by IBM Spectrum Protect and so that the client cannot delete archived files.

```
update node larry compression=yes archdelete=no
```

Example: Update a node's number of files that can be transferred as a group

Update node LARRY and increase the TXNGroupmax value to 1,000.

```
update node larry txngroupmax=1000
```

Example: Update a node and allow it to deduplicate on the client

Update a node BOB so that it can deduplicate on the client.

```
update node bob deduplication=clientorserver
```

Example: Update the role of node BOB to a server-device for PVU estimation reporting

If you want to accumulate PVU values, only server device roles are recorded. You can update a node from client-device to server-device by issuing the UPDATE NODE command. For this example, node BOB is updated to a server-device.

```
update node bob role=server
```

Example: Update a node definition on a source replication server

NODE1 is defined to a source replication server. The data that belongs to NODE1 was previously exported to a target replication server. Update the replication rule for backup data that belongs to NODE1 so that active backup data is replicated with a high priority. Enable replication for the node. Set up data synchronization with the target replication server.

```
update node node1 bkreplruledefault=active_data_high_priority  
replstate=enabled replmode=synccsend
```

Example: Update a node definition to enable recovery of damaged files

Update the PAYROLL node to enable the recovery of damaged files from a target replication server.

```
update node payroll recoverdamaged=yes
```

Related commands

Table 2. Commands related to UPDATE NODE

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PVUESTIMATE	Displays an estimate of the client-devices and server-devices being managed.
QUERY REPLNODE	Displays information about the replication status of a client node.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REMOVE REPLNODE	Removes a node from replication.
RENAME NODE	Changes the name for a client node.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
SET DEDUPVERIFICATIONLEVEL	Specifies the percentage of extents verified by the server during client-side deduplication.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
SET REPLRECOVERDAMAGED	Specifies whether node replication is enabled to recover damaged files from a target replication server.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE FILESPACE	Changes file-space node-replication rules.

Related reference:

[Ssl client option](#)

UPDATE NODEGROUP (Update a node group)

Use this command to modify the description of a node group.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
>>-UPDATE NODEGroup--group_name--DESCRiption---description---<<
```

Parameters

group_name

Specifies the name of the node group whose description you want to update.

DESCRiption (Required)

Specifies a description of the node group. This parameter is required. The maximum length of the description is 255 characters. If the description contains any blanks, enclose the entire description in quotation marks.

Example: Update a node group's description

Update the node group, `group1`, with a new description.

```
update nodegroup group1 description="Human Resources"
```

Related commands

Table 1. Commands related to UPDATE NODEGROUP

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.

UPDATE PATH (Change a path)

Use this command to update a path definition.

Syntax and parameter descriptions are available for the following path types.

- UPDATE PATH (Change a path when the destination is a drive)
- UPDATE PATH (Change a path when the destination is a library)
- **AIX** | **Linux** UPDATE PATH (Update a path when the destination is a ZOSMEDIA library)

For detailed and current device support information, see the Supported Devices website for your operating system:

- **AIX** | **Windows** Supported devices for AIX and Windows
- **Linux** Supported devices for Linux

Related commands

Table 1. Commands related to UPDATE PATH

Command	Description
DEFINE DATAMOVER	Defines a data mover to the IBM Spectrum Protect server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE PATH	Deletes a path from a source to a destination.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DATAMOVER	Changes the definition for a data mover.

UPDATE PATH (Change a path when the destination is a drive)

Use this syntax when updating a path definition to a drive.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate PATH--source_name--destination_name----->
>--SRCType-----+---+DATAMover-+---+-----+----->
           '-SERVer----'   '-AUTODetect-----+No---+'
                                   '-Yes-'
>--DESTType-----Drive--LIBRARY-----library_name----->
>--+-----+-----+-----+----->
   '-DEVIce-----device_name-'   '-ONLine-----+Yes-+-'
                                   '-No--'
>--+-----+-----+-----+----->>
|           .,-----, |
|           v           | |
'-DIRectory-----directory_name-+-'
```

Parameters

source_name (Required)

Specifies the name of source for the path. This parameter is required.

destination_name (Required)

Specifies the name of the destination. This parameter is required.

SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVer

Specifies that a server or a storage agent is the source.

AUTODetect

Specifies whether the serial number for a drive or library will be automatically detected, reported, and updated in IBM Spectrum Protect™. This parameter is optional. This parameter is only valid for paths defined from the local server to a drive or a library. Possible values are:

No

Specifies that the serial number is not automatically updated.

Yes

Specifies that the serial number is automatically updated to reflect the same serial number that the drive reports to IBM Spectrum Protect.

Important:

1. If you have not previously entered a serial number, then AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.
2. AUTODETECT=YES in this command overrides the serial number set in the DEFINE DRIVE command.
3. If you set DESTTYPE=DRIVE and AUTODETECT=YES, then the drive element number in the IBM Spectrum Protect database will be automatically changed to reflect the same element number that corresponds to the serial number of that drive. This is true for drives in a SCSI library. For more information about the element number, see the DEFINE DRIVE command.
4. Depending on the capabilities of the device, the AUTODETECT parameter may not be supported.

DESTType=DRive (Required)

Specifies that a drive is the destination. When the destination is a drive, you must specify a library name. This parameter is required.

LIBRARY

Specifies the name of the library to which the drive is assigned. The library and its drives must already be defined to the server. If the path is from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349x, or ACSLS.

DEVIce

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

AIX The source uses the device name to access the drive. See Table 1 for examples.

Table 1. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	AIX /dev/rmt3
Storage agent to a drive (not a FILE drive)	mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM® System Storage® N Series: rst01

Linux The source uses the device name to access the drive. See Table 2 for examples.

Table 2. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	/dev/tmsmcsi/mt3
Storage agent to a drive (not a FILE drive)	/dev/tmsmcsi/mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM System Storage N Series: rst01

Windows The source uses the device name to access the drive. See Table 3 for examples.

Table 3. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	Windows mt3
Server to a drive (REMOVABLEFILE drive)	e:
Storage agent to a drive (not a FILE drive)	mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM System Storage N Series: rst01

Important:

- For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. For more information, see *IBM Tape Device Drivers Installation and User's Guide*, which can be downloaded from the IBM Systems support site at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the SYSCONFIG command. Use this command to determine device names for drives:

```
sysconfig -t
```

ONLine

Specifies whether the path is available for use. This parameter is optional. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

For example, if the path from a data mover to a drive is online, but either the data mover or the drive is offline, you cannot use the path.

DIRectory

Specifies the directory location or locations for a storage agent to access the files in a FILE library. The DIRECTORY parameter is also used for devices of type REMOVABLEFILE. For REMOVABLEFILE devices, the DIRECTORY parameter provides information for the server (not a storage agent) along with the DRIVE parameter to describe access to the device. This parameter is optional.

On storage agents, this parameter is only valid when *all* of the following conditions are true:

- The source type is SERVER (meaning a storage agent that has been defined as a server to this server).
- The source name is the name of a storage agent, *not* the server.
- The destination is a logical drive that is part of a FILE library.
- If multiple directories were specified for the device class associated with the FILE library, the same number of directories must be specified with the DIRectory parameter of the DEFINE PATH command, for each drive in the FILE library. Storage agent directories are not validated on the server. Specifying incorrect directories can cause a run-time failure.

The directory name or names identify the locations where the storage agent reads and writes the files that represent storage volumes for the FILE device class that is associated with the FILE library. The default value for DIRECTORY is the directory of the server at the time the command is issued.

Use a naming convention that you can use to associate the directory with a particular physical drive. This can help ensure that your configuration is valid for sharing the FILE library between the server and storage agent. If the storage agent is on a Windows system, use a universal naming convention (UNC) name. When the storage agent lacks permission to access remote storage, the storage agent will experience mount failures.

Windows The account associated with the storage agent service must be either an account within the local administrator's group or an account within the domain administrator's group. If the account is in the local administrator's group, the user ID and password must match that of an account with permissions to access storage as provided by the machine which administers the remote share. For example, if a SAMBA server is providing access to remote storage, the user ID and password in the SAMBA configuration must match that of the local administrator user ID and password associated with the storage agent service.

```
define devclass file devtype=file shared=yes mountlimit=1
directory=d:\filedir\dir1
define path stal file1 srctype=server desttype=drive
library=file1 device=file directory=\\192.168.1.10\filedir\dir1
```

In the previous example, the DEFINE DEVCLASS command establishes the shared file system in the directory accessed by the server as D:\FILEDIR\DIR1. The storage agent, however, is using UNC name \\192.168.1.10\FILEDIR\DIR1. This means that the machine with TCP/IP address 192.168.1.10 is sharing the same directory using FILEDIR as the shared name. Also, the storage agent service has an account which can access this storage. It can access it either because it is associated with a local account with the same user ID and password as 192.168.1.10 or it is associated with a domain account which is available on both the storage agent and on 192.168.1.10. If appropriate to the installation, you can replace the 192.168.1.10 with a symbolic name such as:

example.yourcompany.com

Important:

- IBM Spectrum Protect does not create shares or permissions, or mount the target file system. You must perform these actions before starting the storage agent.
- You can modify a list of directories only by replacing the entire list.
- You must ensure that storage agents can access newly created FILE volumes. To access FILE volumes, storage agents replace names from the directory list in the device-class definition with the names in the directory list for the associated path definition. The following illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created FILE volumes.

Suppose you want to use these three directories for a FILE library: **Windows**

- c:\server
 - d:\server
 - e:\server
- | | |
|------------|--------------|
| AIX | Linux |
|------------|--------------|
- /opt/tivoli1

- o /opt/tivoli2
 - o /opt/tivoli3
1. You use the following command to set up a FILE library named CLASSA with one drive named CLASSA1 on *SERVER1*: **Windows**

```
define devclass classa devtype=file
directory="c:\server,d:\server,e:\server"
shared=yes mountlimit=1
```

AIX | Linux

```
define devclass classa devtype=file
directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
shared=yes mountlimit=1
```

2. You want the storage agent *STA1* to be able to use the FILE library, so you define the following path for storage agent *STA1*: **Windows**

```
define path server1 sta1 srctype=server desttype=drive device=file
directory="\\192.168.1.10\c\server,\\192.168.1.10\d\server,
\\192.168.1.10\e\server" library=classa
```

Windows In this scenario, the storage agent, *STA1*, will replace the directory name *c:\server* with the directory name *\\192.168.1.10\c\server* to access FILE volumes that are in the *c:\server* directory on the server.

AIX | Linux

```
define path server1 sta1 srctype=server desttype=drive device=file
directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

AIX | Linux In this scenario, the storage agent, *STA1*, will replace the directory name */opt/tivoli1* with the directory name */opt/ibm1/* to access FILE volumes that are in the */opt/tivoli1* directory on the server.

3. **Windows** File volume *c:\server\file1.dsm* is created by *SERVER1*. If you later change the first directory for the device class with the following command:

```
update devclass classa directory="c:\otherdir,d:\server,e:\server"
```

SERVER1 will still be able to access file volume *c:\server\file1.dsm*, but the storage agent *STA1* will not be able to access it because a matching directory name in the *PATH* directory list no longer exists. If a directory name is not available in the directory list associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume will still be accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

4. **AIX | Linux** If file volume */opt/tivoli1/file1.dsm* is created on *SERVER1*, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,
/opt/tivoli3"
```

SERVER1 will still be able to access file volume */opt/tivoli1/file1.dsm*, but the storage agent *STA1* will not be able to access it because a matching directory name in the *PATH* directory list no longer exists. If a directory name is not available in the directory list associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume will still be accessible from the server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

Example: Update a path from a data mover NAS file server to a tape drive

Update a path from a data mover that is a NAS file server to the drive *TAPEDRV2* that the data mover uses for backup and restore operations. In this example, the NAS data mover is *NAS1*, the library is *NASLIB*, and the device name for the drive is *rst01*.

```
update path nas1 tapedrv2 srctype=datamover desttype=drive library=naslib
device=rst01
```

UPDATE PATH (Change a path when the destination is a library)

Use this syntax when updating a path definition to a library.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate PATH--source_name--destination_name----->
>--SRCType-----+--DATAMover-+--+-----+----->
                '-SERVer----'   '-AUTODetect-----+No--+-'
                                     '-Yes-'
>--DESTType-----LIBRary-+-----+----->
                +-DEVIce-----device_name-----+
                '-EXTERNALManager---path_name-'
>--+-----+-----><
  '-ONLine-----+Yes-+-'
                '-No--'
```

Parameters

source_name (Required)

Specifies the name of source for the path. This parameter is required.

destination_name (Required)

Specifies the name of the destination. This parameter is required.

Important: To define a path from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349X, or Automated Cartridge System Library Software (ACSL).

SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVer

Specifies that a server or a storage agent is the source.

AUTODetect

Specifies whether the serial number for a drive or library is automatically detected, reported, and updated in IBM Spectrum Protect™. This parameter is optional. This parameter is only valid for paths defined from the local server to a library.

Possible values are:

No

Specifies that the serial number is not automatically updated.

Yes

Specifies that the serial number is automatically updated to reflect the same serial number that the drive reports to IBM Spectrum Protect.

Important:

1. If you have not previously entered a serial number, then AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.
2. AUTODETECT=YES in this command overrides the serial number set in the DEFINE DRIVE command.
3. Depending on the capabilities of the device, the AUTODETECT parameter may not be supported.

DESTType=LIBRary (Required)

Specifies that a library is the destination.. This parameter is required.

DEVIce

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

AIX The source uses the device name to access the drive or library. See Table 1 for examples.

Table 1. Examples of device names

Source to destination	Example
-----------------------	---------

Source to destination	Example
Server to a library	AIX /dev/lb4 Linux /dev/tmsmcsi/lb4
NAS data mover to a library	mc0

Linux The source uses the device name to access the drive or library. See Table 2 for examples.

Table 2. Examples of device names

Source to destination	Example
Server to a library	/dev/tmsmcsi/lb4
NAS data mover to a library	mc0

Windows The source uses the device name to access the drive or library. See Table 3 for examples.

Table 3. Examples of device names

Source to destination	Example
Server to a library	Windows lb4.1
NAS data mover to a library	mc0

Important:

- For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. For more information, see *IBM Tape Device Drivers Installation and User's Guide*, which can be downloaded from the IBM® Systems support site at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the SYSCONFIG command. Use this command to determine the device name for a library:

```
sysconfig -m
```

EXTERNALManager

Specifies the location of the external library manager where IBM Spectrum Protect can send media access requests. Use single quotation marks around the value of this parameter. For example, enter: **AIX**

```
/usr/lpp/GESedt-acsls/bin/elmdt
```

Linux

```
/opt/GESedt-acsls/bin/elmdt
```

Windows

```
C:\Program Files\GES\EDT-ACSLs\bin\elmdt.exe
```

This parameter is required when the library name is an external library.

ONLine

Specifies whether the path is available for use. This parameter is optional. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Important: If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized.

AIX

Linux

UPDATE PATH (Update a path when the destination is a ZOSMEDIA library)

Use this syntax when you update a path to a ZOSMEDIA library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate PATH--source_name--destination_name----->
>--SRCType-----SERVer--DESTType-----LIBRary----->
>--ZOSMEDIASERVER-----server_name--+-----+-----<
                                '-ONLine-----+Yes--+'
                                '-No--'
```

Parameters

source_name (Required)

Specifies the name of source for the path.

destination_name (Required)

Specifies the name of the destination.

SRCType=SERVer (Required)

Specifies that the IBM Spectrum Protect™ server or a storage agent is the source.

DESTType=LIBRary (Required)

Specifies that a library is the destination.

ZOSMEDIAServer (Required)

Specifies the server name that represents a Tivoli® Storage Manager for z/OS® Media server.

ONLine

Specifies whether the path is available for use. This parameter is optional. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Important: If the path to a library is offline, the server cannot access the library. If the server is halted and restarted while the path to the library is offline, the library is not initialized during server initialization. The path must be updated to ONLINE=YES to access the library.

UPDATE POLICYSET (Update a policy set description)

Use this command to change the description of a policy set. You cannot change the description of the ACTIVE policy set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-UPDate Policyset--domain_name--policy_set_name----->
>--DESCRiption-----description-----<
```

Parameters

domain_name (Required)

Specifies the policy domain to which the policy set belongs.

policy_set_name (Required)

Specifies the policy set to update. You cannot change the ACTIVE policy set.

DEscription (Required)

Specifies text that describes the policy set. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("").

Example: Update a policy set

Update a policy set called VACATION for the EMPLOYEE_RECORDS policy domain with a description of "Schedule Planning Information."

```
update policysset employee_records vacation
description="schedule planning information"
```

Related commands

Table 1. Commands related to UPDATE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE MGMTCLASS	Defines a management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

UPDATE PROFILE (Update a profile description)

Use this command on a configuration manager to update a profile description.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate PROFile--profile_name--DEscription---description---<<
```

Parameters

profile_name (Required)

Specifies the profile to update.

DEscription (Required)

Specifies a description for the profile. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a description, specify a null string ("").

Example: Update a profile's description

Update the description for profile DELTA.

```
update profile delta description="PAYROLL domain"
```

Related commands

Table 1. Commands related to UPDATE PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.

UPDATE RECOVERYMEDIA (Update recovery media)

Use this command to update information about recovery media.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate RECOVERYMedia--media_name----->
>--+-----+-----+----->
|           .-,-----.|
|           v           ||
'-VOLumenames-----volume_name-+-'
>--+-----+-----+-----+----->
'-DESCRiption----description-' '-LOcation----location-'
>--+-----+-----+-----+----->
'-Type-----+B0ot--+-' '-PR0duct----product_name-'
'-Other-'
>--+-----+-----+----->>
'-PR0DUCTInfo----product_information-'
```

Parameters

media_name (Required)

Specifies the name of the recovery media to be updated.

VOLumenames

Specifies the names of volumes that contain the recoverable data (for example, operating system image copies). If you specify a TYPE=BOOT, you must specify the boot media volume names in the order in which they are to be loaded at recovery time. The volume names list can be up to 255 characters. Enclose the list in quotation marks if it contains any blank characters. To remove all volume names, specify a null string ("").

DESCRiption

Specifies the description of the recovery media. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters.

LOcation

Describes the location of the recovery media. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove a location description, specify a null string ("") for the value.

Type

Specifies the type of recovery media. This parameter is optional. Possible values are:

BOot

Specifies that this is boot media. You must specify volume names if the type is BOOT.

OTHer

Specifies that this is not boot media. For example, a CD that contains operating system manuals.

PROduct

Specifies the name of the product that wrote to this media. This parameter is optional. You can use up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove a product name, specify a null string ("") for the value.

PRODUCTInfo

Specifies any information about the product that wrote to the media that you may need to restore the machine. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove previously defined product information, specify a null string ("") for the value.

Example: Update a recovery media's location description

Update the location description for recovery media DIST5RM to "Corporate Headquarters Data Vault."

```
update recoverymedia dist5rm
location="Corporate Headquarters Data Vault"
```

Related commands

Table 1. Commands related to UPDATE RECOVERYMEDIA

Command	Description
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.

UPDATE REPLRULE (Update replication rules)

Use this command to enable or disable a replication rule.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate REPLRule--rule_name----StAtE-----+ENabled--+------><
                                     '-DISabled-'
```

Parameters

rule_name (Required)

Specifies the name of the replication rule to be updated. You can use wildcard characters to specify one or more rules. You can specify one of the following rules:

- ALL_DATA
- ACTIVE_DATA

- ALL_DATA_HIGH_PRIORITY
- ACTIVE_DATA_HIGH_PRIORITY

State (Required)

Specifies whether replication is allowed for the rule. You can specify one of the following values:

Enabled

Specifies that the data to which the rule applies is ready to be replicated

Disabled

Specifies that replication does not occur until you enable it.

Example: Disable replication for backup data

Disable replication of normal-priority, active-backup data for all file spaces in all client nodes that are configured for replication:

```
update replrule active_data state=disabled
```

Related commands

Table 1. Commands related to UPDATE REPLRULE

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes that are associated with a client node.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

UPDATE SCHEDULE (Update a schedule)

Use this command to update a client or administrative command schedule.

The UPDATE SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. Within these two forms, you can select either classic or enhanced style schedules. The syntax and parameters for each form are defined separately.

Table 1. Commands related to UPDATE SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY SCHEDULE	Displays information about schedules.


```

>--+-----+-----+-----+----->
  '-STARTDate---date-' '-STARTTime---time-'

>--+-----+-----+-----+----->
  '-DURation---number-' '-DURUnits---+Minutes---+'
                                     +-Hours-----+
                                     +-Days-----+
                                     '-INDefinite-'

>--+-----+-----+-----+----->
  '-MAXRUNTime---number-' '-SCHEDStyle---Classic-'

>--+-----+-----+-----+----->
  '-PERiod---number-' '-PERUnits---+Hours---+'
                                     +-Days---+
                                     +-Weeks---+
                                     +-Months---+
                                     +-Years---+
                                     '-Onetime-'

>--+-----+-----+-----+----->
  '-DAYofweek---+ANY---+'
      +-WEEKDay---+
      +-WEEKEnd---+
      +-SUnDay---+
      +-MonDay---+
      +-TUESday---+
      +-WednesDay+
      +-THURsday--+
      +-FRIday---+
      '-SATurday--'

>--+-----+-----+-----+-----><
  '-EXPIration---+Never---+'
      '-date--'

```

Notes:

1. You must specify at least one optional parameter on this command.

Syntax for an enhanced client schedule

```

(1)
>>-UPDate SChedule-----domain_name--schedule_name----->

>--+-----+-----+-----+----->
  '-Type---Client-' '-DEScRiption---description-'

>--+-----+-----+-----+----->
  '-ACTion---+Incremental-----+'
      +-Selective-----+
      +-Archive---+-----+
      |           '-SUBACTion---+-----+' |
      |                                     '-FASTBack-' |
      +-Backup---+-----+
      |           '-SUBACTion---+-----+' |
      |                                     +-FASTBack---+ |
      |                                     +-SYSTEMState+ |
      |                                     +-VApp-----+ |
      |                                     '-VM-----+' |
      +-REStore-----+
      +-RETRieve-----+
      +-IMAGEBACkup-----+
      +-IMAGERESTore-----+
      +-Command-----+
      '-Macro-----+'

>--+-----+-----+-----+----->
  '-OPTions---option_string-'

>--+-----+-----+-----+----->

```

```

'-OBJects---object_string-' '-PRIority---number-'
>----->
'-STARTDate---date-' '-STARTTime---time-'
>----->
'-DURation---number-' '-DURUnits---Minutes--'
                                     +-Hours---+
                                     '-Days----'
>----->
'-MAXRUNtime---number-' '-SCHEDStyle---Enhanced-'
>----->
'-MONth---ANY-----+' '-DAYOFMonth---ANY--+'
      +-JANuary---+
      +-February--+
      +-MARch-----+
      +-April-----+
      +-May-----+
      +-JUNe-----+
      +-JULy-----+
      +-AUGust----+
      +-September--+
      +-October---+
      +-November--+
      '-December--'
>----->
'-WEEKofmonth---ANY---+'
      +-First--+
      +-Second+
      +-Third--+
      +-FOurth+
      '-Last---'
>----->
'-DAYofweek---ANY-----+'
      +-WEEKDay---+
      +-WEEKEnd---+
      +-SUNday---+
      +-Monday----+
      +-TUESday---+
      +-WEdnesday+
      +-THursday--+
      +-FRIday---+
      '-SATurday--'
>----->
'-EXPIration---Never--+
      '-date--'

```

Notes:

1. You must specify at least one optional parameter on this command.

Parameters

domain_name (Required)

Specifies the name of the policy domain to which this schedule belongs.

schedule_name (Required)

Specifies the name of the schedule to be updated.

Type=Client

Specifies that a client schedule is updated. This parameter is optional. The default is CLIENT.

DESCRIPTION

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters. To remove a previously defined description, specify a null string ("") for this value.

ACTION

Specifies the action that occurs when this schedule is processed. Possible values are:

Incremental

Specifies that the schedule backs up all files that are new or that have changed since the last incremental backup. Incremental also backs up any file for which all existing backups might have expired.

Selective

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

Archive

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

Backup

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

REStore

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

REtrieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

Remember: A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.

IMAGERESTore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.

Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

SUBACTion

You can specify one of the following values:

""

When a null string (two double quotes) is specified with ACTION=BACKUP the backup is an incremental.

FASTBBack

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

SYSTEMState

Specifies that a client Systemstate backup is scheduled.

VApp

Specifies that a client vApp backup is scheduled. A vApp is a collection of pre-deployed virtual machines.

VM

Specifies that a client VMware backup operation is scheduled.

Deploy

Specifies whether to update client workstations with deployment packages that are specified with the OBJECTS parameter. The OBJECTS parameter must contain two specifications, the package files to retrieve and the location from which to retrieve them. Ensure that the objects are in the order *files location*. For example:

```
define schedule standard deploy_1 action=DEPLOY objects=  
"\\IBM_ANR_WIN\c$\tsm\maintenance\client\v6r2\Windows\X32\v620\v6200\  
..\IBM_ANR_WIN\"
```

Values for the following options are restricted when you specify ACTION=DEPLOY:

PERUNITS

Specify PERUNITS=ONETIME. If you specify PERUNITS=PERIOD, the parameter is ignored.

DURUNITS

Specify MINUTES, HOURS, or DAYS for the DURUNITS parameter. Do not specify INDEFINITE.

SCHEDSTYLE

Specify the default style, CLASSIC.

The SCHEDULE command fails if the parameters do not conform to the required parameter values, such as the V.R.M.F.

OPTions

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no effect when the client processes the scheduled command:

- MAXCMDRETRIES
- OPTFILE
- QUERYSCHEDPERIOD
- RETRYPERIOD
- SCHEDLOGNAME
- SCHEDMODE
- SERVERNAME
- TCPCLIENTADDRESS
- TCPCLIENTPORT

Windows When you define a scheduler service by using the DSMCUTIL command or the backup-archive client GUI wizard, you specify an options file. You cannot override the options in that options file by issuing the scheduled command. You must modify the options in your scheduler service.

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify `subdir=yes` and `domain all-local -systemobject`, enter:
 - `options='-subdir=yes -domain="all-local -c: -systemobject"'`
- To specify `domain all-local -c: -d:`, enter:
 - `options='-domain="all-local -c: -d:"'`

Windows Tip:

For Windows clients running in batch mode, if the use of quotation marks is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

OBJects

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when ACTION=INCREMENTAL. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify ACTION=INCREMENTAL without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:

```
ANS1082E Invalid destination file specification <filespec> entered.
```

- If you specify more than two file specifications, you receive this error:

```
ANS1102E Excessive number of command line arguments passed to the program!
```

When you specify ACTION=ARCHIVE, INCREMENTAL, or SELECTIVE for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

Windows If you are using characters that have a special meaning for Windows users, such as commas, surround the entire argument in two pairs of double quotes, then surround the entire string with single quotes. The following examples show you how to specify some file names:

- To specify C:\FILE 2, D:\GIF FILES, and E:\MY TEST FILE, enter:
 - OBJECTS="C:\FILE 2" "D:\GIF FILES" "E:\MY TEST FILE"
- To specify D:\TEST FILE, enter:
 - OBJECTS="D:\TEST FILE"
- To specify D:TEST,FILE:
 - OBJECTS="\"D:\TEST, FILE\""

AIX **Linux** The following examples show how to specify some file names:

- To specify /home/file 2, /home/gif files, and /home/my test file, enter:
 - OBJECTS="/home/file 2" "/home/gif files" "/home/my test file"
- To specify /home/test file, enter:
 - OBJECTS="/home/test file"

Windows Tip:

For Windows clients running in batch mode, if the use of double quotes is necessary, use interactive mode or operating system escape characters. For additional information, see the following topics:

- Processing a series of commands from the administrative client
- Processing individual commands from the administrative client

PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when IBM Spectrum Protect™ processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the STARTTIME parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00. If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00.

DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

Tip: Define schedules with durations longer than 10 minutes. Doing this will give the IBM Spectrum Protect scheduler enough time to process the schedule and prompt the client.

DURUnits

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is HOURS.

Use this parameter with the DURATION parameter to specify how long the startup window remains open to process the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. You can specify one of the following values:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE, unless you specify PERUNITS=ONETIME. The INDEFINITE value is not allowed with enhanced schedules.

MAXRUNtime

Specifies the maximum run time, which is the number of minutes during which all client sessions that are started by the scheduled operation should be completed. If sessions are still running after the maximum run time, the server issues a warning message, but the sessions continue to run.

Tip: The maximum run time is calculated from the beginning of the startup window and not from the time that sessions start within the startup window.

Restrictions:

- The value of the parameter is not distributed to servers that are managed by an enterprise configuration manager.
- The value of the parameter is not exported by the EXPORT command.

The parameter is optional. You can specify a number in the range 0-1440. A value of 0 means that the maximum run time is indefinite, and no warning message is issued. The maximum run time must be greater than the startup window duration, which is defined by the DURATION and DURUNITS parameters.

For example, if the start time of a scheduled operation is 9:00 PM, and the duration of the startup window is 2 hours, the startup window is 9:00 PM - 11:00 PM. If the maximum run time is 240 minutes, that is, 4 hours, all client sessions for this operation should be completed by 1:00 AM. If one or more sessions are still running after 1:00 AM, the server issues a warning message.

Tip: Alternatively, you can specify a *run time alert* value of 1:00 AM in the IBM Spectrum Protect Operations Center.

SCHEDStyle

This parameter is optional. SCHEDSTYLE defines either the interval between times when a schedule can run, or the days on which it can run. The style can be either classic or enhanced. This parameter must be specified when you change a schedule from classic to enhanced or back to classic. Otherwise, the value for the existing schedule is used.

For classic schedules, these parameters are allowed: PERIOD, PERUNITS, and DAYOFWEEK. These parameters are not allowed: MONTH, DAYOFMONTH, and WEEKOFMONTH. If the previous schedule style was enhanced, the MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK parameters are reset. DAYOFWEEK, PERIOD, and PERUNITS are set to default values unless they are specified with the update command.

For enhanced schedules, these parameters are allowed: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. These parameters are not allowed: PERIOD and PERUNITS. If the previous schedule style was classic, the DAYOFWEEK, PERIOD, and PERUNITS parameters are reset. MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK are set to default values unless they are specified with the update command.

PERiod

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is DAYS.

Use this parameter with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. You can specify one of the following values:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify PERUNITS=MONTHS, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be

processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

When you specify PERUNITS=YEARS, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the PERIOD parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the DAYofweek parameter, depending on whether the schedule style was defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. DAYOFWEEK must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

Possible values for the DAYofweek parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUnday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

TUesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

THursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

Saturday

Specifies that the startup window begins on Saturday.

MONTH

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY, which means that the schedule runs during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs on each of the specified days of the month. If multiple values resolve to the same day, the schedule runs only once that day.

The default value is ANY, which means that the schedule runs on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

If an existing schedule specifies a value other than ANY for DAYOFWEEK and WEEKOFMONTH, and DAYOFMONTH is updated, DAYOFWEEK and WEEKOFMONTH are reset to ANY.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule runs only once during that week.

The default value is ANY. ANY means that the schedule runs during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXpiration

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. You can specify one of the following values:

Never

Specifies that the schedule never expires.

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Update the priority of a schedule

Update the MONTHLY_BACKUP schedule that belongs to the STANDARD policy domain by setting its priority value to 1.

```
update schedule standard monthly_backup priority=1
```

Example: Update the expiration date of a schedule

Update the WEEKLY_BACKUP schedule that belongs to the EMPLOYEE_RECORDS policy domain to expire on March 29, 1999 (03/29/1999).

```
update schedule employee_records weekly_backup expiration=03/29/1999
```

Example: Update a schedule to archive on the last Friday of a month

Update a schedule from archiving files quarterly on the last Friday of the month to archiving on the last day of the specified months.

```
update schedule employee_records quarterly_archive dayofmonth=-1
```

WEEKOFMONTH and DAYOFWEEK are reset to ANY.

UPDATE SCHEDULE (Update an administrative schedule)

Use this command to update selected parameters for an administrative command schedule.

You cannot schedule MACRO or QUERY ACTLOG commands.

A managed administrative schedule that is updated by a configuration manager is set to an inactive state on the managed servers during configuration refresh processing. It remains in an inactive state until it is updated to an active state on those servers.

Privilege class

To update an administrative schedule, you must have system privilege.

Syntax

Classic administrative schedule

```
(1)
>>-UPDate SChedule-----schedule_name----->
>--+-----+-----+-----+----->
  '-Type-----Administrative-'  '-CMD-----command-'
>--+-----+-----+-----+----->
  '-ACTIVE-----+Yes+-'  '-DEScRiption-----description-'
                    '-No--'
>--+-----+-----+-----+----->
  '-PRIority-----number-'  '-STARTDate-----date-'
>--+-----+-----+-----+----->
  '-STARTTime-----time-'  '-DURation-----number-'
>--+-----+-----+-----+----->
  '-DURUnits-----+Minutes-----+'  '-MAXRUNTime-----number-'
                    +-Hours-----+
                    +-Days-----+
                    '-INDefinite-'
>--+-----+-----+-----+----->
  '-SCHEdStyle-----Classic-'  '-PERiod-----number-'
>--+-----+-----+-----+----->
  '-PERUnits-----+Hours----+'
                    +-Days----+
                    +-Weeks---+
                    +-Months--+
                    +-Years---+
                    '-Onetime-'
>--+-----+-----+-----+----->
  '-DAYofweek-----+ANY-----+'
                    +-WEEKDay---+
                    +-WEEKEnd---+
                    +-Sunday----+
                    +-Monday----+
                    +-TUESday---+
                    +-WednesDay+
                    +-THursDay--+
                    +-Friday----+
                    '-SATurday--'
>--+-----+-----+-----+----->>
  '-EXPIration-----+Never+-'
                    '-date--'
```

Notes:

1. You must specify at least one optional parameter on this command.

Syntax

Enhanced administrative schedule

```
(1)
>>-UPDate SChedule-----schedule_name----->
>--+-----+-----+-----+----->
>  '-Type-----Administrative-'  '-CMD-----command-'
>--+-----+-----+-----+----->
>  '-ACTIVE-----+Yes-+-'  '-DESCRiption-----description-'
>    '-No--'
>--+-----+-----+-----+----->
>  '-PRIority-----number-'  '-STARTDate-----date-'
>--+-----+-----+-----+----->
>  '-STARTTime-----time-'  '-DURation-----number-'
>--+-----+-----+-----+----->
>  '-DURUnits-----+Minutes-+-'  '-MAXRUNtime-----number-'
>    +-Hours---+
>    '-Days----'
>--+-----+-----+-----+----->
>  '-SCHEDStyle-----Enhanced-'  '-MONth-----+ANY-----+-'
>    +-JAnuary---+
>    +-FebruAry--+
>    +-MARCh-----+
>    +-APRil-----+
>    +-May-----+
>    +-JUNe-----+
>    +-JULy-----+
>    +-AUGust----+
>    +-SeptemBer--+
>    +-October---+
>    +-November--+
>    '-December--'
>--+-----+-----+-----+----->
>  '-DAYOFMonth-----+ANY-+-'  '-WEEKofmonth-----+ANY-----+-'
>    '-Day-'  +-First--+
>    +-Second--+
>    +-Third--+
>    +-FOurth--+
>    '-Last---'
>--+-----+-----+-----+----->
>  '-DAYofweek-----+ANY-----+-'
>    +-WEEKDay---+
>    +-WEEKEnd---+
>    +-SUNday----+
>    +-MONday----+
>    +-TUESday---+
>    +-WEDnesday--+
>    +-THURsday--+
>    +-FRIday----+
>    '-SATurday--'
>--+-----+-----+-----+----->>
>  '-EXPIration-----+Never-+-'
>    '-date--'
```

Notes:

1. You must specify at least one optional parameter on this command.

Parameters

schedule_name (Required)

Specifies the name of the schedule to be updated.

Type=Administrative (Required)

Specifies that an administrative command schedule is updated.

CMD

Specifies the administrative command to be scheduled for processing. This parameter is optional. The command you specify can contain up to 512 characters. Enclose the command in quotation marks if it contains blanks.

You cannot specify redirection characters with this parameter.

ACTIVE

Specifies whether the administrative command is eligible for processing. This parameter is optional. An administrative command schedule will not be processed unless it is set to the active state. Possible values are:

YES

Specifies that the administrative command is eligible for processing.

NO

Specifies that the administrative command is not eligible for processing.

DEScRiption

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blanks. To remove a previously defined description, specify a null string ("") for this value.

PRIOrity

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when IBM Spectrum Protect™ processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the STARTTIME parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or +days	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08

Value	Description	Example
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00. If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00.
NOW-HH:MM or - HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00.

DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

DURUnits

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is HOURS.

Use this parameter with the DURATION parameter to specify how long the startup window remains open to process the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. You can specify one of the following values:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE, unless you specify PERUNITS=ONETIME. The INDEFINITE value is not allowed with enhanced schedules.

MAXRUNtime

Specifies the maximum run time, which is the number of minutes during which server processes that are started by the scheduled commands must be completed. If processes are still running after the maximum run time, the central scheduler cancels the processes.

Tips:

- The processes might not end immediately when the central scheduler cancels them; they end when they register the cancellation notification from the central scheduler.
- The maximum run time is calculated beginning from when the server process starts. If the schedule command starts more than one process, each process maximum run time is calculated from when the process starts.
- This parameter does not apply to some processes, such as duplicate-identification processes, which can continue to run after the maximum run time.
- This parameter does not apply if the scheduled command does not start a server process.
- Another cancel time might be associated with some commands. For example, the MIGRATE STGPOOL command can include a parameter that specifies the length of time that the storage pool migration runs before the migration is

automatically canceled. If you schedule a command for which a cancel time is defined, and you also define a maximum run time for the schedule, the processes are canceled at whichever cancel time is reached first.

Restrictions:

- The value of the parameter is not distributed to servers that are managed by an enterprise configuration manager.
- The value of the parameter is not exported by the EXPORT command.

This parameter is optional. You can specify a number in the range 0-1440. A value of 0 means that the maximum run time is indefinite, and the central scheduler does not cancel processes. The maximum run time must be greater than the startup window duration, which is defined by the DURATION and DURUNITS parameters.

For example, if the start time of a scheduled command is 9:00 PM, and the duration of the startup window is 2 hours, the startup window is 9:00 PM - 11:00 PM. If the maximum run time is 240 minutes, that is, 4 hours, all applicable server processes that are started by the command must be completed by 1:00 AM. If one or more applicable processes are still running after 1:00 AM, the central scheduler cancels the processes.

Tip: Alternatively, you can specify an *end time* of 1:00 AM in the IBM Spectrum Protect Operations Center.

SCHEDStyle

This parameter is optional. SCHEDSTYLE defines either the interval between times when a schedule should run, or the days on which it should run. The style can be either classic or enhanced. This parameter must be specified when you change a schedule from classic to enhanced or back to classic. Otherwise, the value for the existing schedule is used.

For classic schedules, these parameters are allowed: PERIOD, PERUNITS, and DAYOFWEEK. These parameters are not allowed: MONTH, DAYOFMONTH, and WEEKOFMONTH. If the previous schedule style was enhanced, the MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK parameters will be reset. DAYOFWEEK, PERIOD, and PERUNITS will be set to default values unless they are specified with the update command.

For enhanced schedules, these parameters are allowed: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. These parameters are not allowed: PERIOD and PERUNITS. If the previous schedule style was classic, the DAYOFWEEK, PERIOD, and PERUNITS parameters will be reset. MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK will be set to default values unless they are specified with the update command.

PERiod

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is DAYS.

Use this parameter with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. You can specify one of the following values:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify PERUNITS=MONTHS, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter,

all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

When you specify PERUNITS=YEAR, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the PERIOD parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the DAYofweek parameter, depending on whether the schedule style was defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. DAYOFWEEK must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

Possible values for the DAYofweek parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

Sunday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

Tuesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAaturday

Specifies that the startup window begins on Saturday.

MONTH

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY. This means the schedule will run during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter can only be specified with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, etc. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run on each of the specified days of the month. If multiple values resolve to the same day, the schedule will run only once that day.

The default value is ANY. This means the schedule will run on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter can only be specified with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule will run only once during that week.

The default value is ANY, meaning the schedule will run during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXpiration

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. You can specify one of the following values:

Never

Specifies that the schedule never expires.

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Update a backup schedule to every three days

Update existing administrative schedule named BACKUP_BACKUPPOOL so that starting today, the BACKUPPOOL primary storage pool is backed up to the COPYSTG copy storage pool every three days at 10:00 p.m.

```
update schedule backup_backuppool type=administrative cmd="backup stgpool
  backuppool copystg" active=yes starttime=22:00 period=3
```

Example: Update a backup schedule to every first and third Friday

Update a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. The existing schedule runs on the first and tenth day of every month. Update it to run the first and third Friday of every month.

```
update schedule backup_archivepool
dayofweek=friday weekofmonth=first,third
```

DAYOFMONTH will be reset to ANY.

UPDATE SCRATCHPADENTRY (Update a scratch pad entry)

Use this command to update data on a line in the scratch pad.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate SCRATCHPadentry--major_category--minor_category----->
```

```
>--subject--Line-----number--Data---data-----><
```

Parameters

major_category (Required)

Specifies the major category in which data is to be updated. This parameter is case sensitive.

minor_category (Required)

Specifies the minor category in which data is to be updated. This parameter is case sensitive.

subject (Required)

Specifies the subject under which data is to be updated. This parameter is case sensitive.

Line (Required)

Specifies the number of the line on which data is to be updated.

Data (Required)

Specifies the new data to be stored on the line. Previous data is deleted. You can enter up to 1000 characters. Enclose the data in quotation marks if the data contains one or more blanks. The data is case sensitive.

Example: Update a scratch pad entry

Update the vacation contact details of an administrator, Jane, in a database that stores information about the location of all administrators:

```
update scratchpentry admin_info location jane line=2 data=
"Out of the office until 18 Nov."
```

Related commands

Table 1. Commands related to UPDATE SCRATCHPADENTRY

Command	Description
DEFINE SCRATCHPADENTRY	Creates a line of data in the scratch pad.
DELETE SCRATCHPADENTRY	Deletes a line of data from the scratch pad.
QUERY SCRATCHPADENTRY	Displays information that is contained in the scratch pad.
SET SCRATCHPADRETENTION	Specifies the amount of time for which scratch pad entries are retained.

UPDATE SCRIPT (Update an IBM Spectrum Protect script)

Use this command to change a command line or to add a new command line to an IBM Spectrum Protect™ script.

Restriction: You cannot redirect the output of a command within an IBM Spectrum Protect script. Instead, run the script and then specify command redirection. For example, to direct the output of script1 to the c:\temp\test.out directory, run the script and specify command redirection as in the following example:

```
run script1 > c:\temp\test.out
```

Privilege class

To issue this command, the administrator must have previously defined the script or must have system privilege.

Syntax

```
>>-UPDate SCRipt--script_name----->
>--+-----+----->
  '-command_line--+-----+'
                    '-Line---number-'
>--+-----+-----><
```

'-DESCRiption-----description-'

Parameters

script_name (Required)

Specifies the name of the script to be updated.

command_line

Specifies a new or updated command to be processed in a script. You must update a command, a description, or both when you issue this command.

Command can contain substitution variables and may be continued across multiple lines if you specify a continuation character (-) as the last character in the command. You can specify up to 1200 characters for the command. Enclose the command in quotation marks if it contains blanks. If you specify this parameter, you can optionally specify the following parameter.

You have the options of running commands serially, in parallel, or serially and in parallel by specifying the SERIAL or PARALLEL script commands for this parameter. You can run multiple commands in parallel and wait for them to complete before proceeding to the next command. Commands will run serially until the parallel command is encountered.

Conditional logic flow statements can be used. These statements include IF, EXIT, and GOTO.

Line

Specifies the line number for the command. If you do not specify a line number, the command line is appended to the existing series of command lines. The appended command line is assigned a line number of five greater than the last command line number in the sequence. For example, if the last line in your script is 015, the appended command line is assigned a line number of 020.

If you specify a line number, the command will replace an existing line (if the number is the same as an existing line). Or the command will insert the specified line (if the line number does not correspond to an existing line number for the command line sequence).

DESCRiption

Specifies a description for the script. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters.

Example: Add a command to the end of a script

Assume that you have defined the following three line script, named QSAMPLE, and that you want to add the QUERY SESSION command to the end of the script.

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS
```

```
update script qsample "query session"
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS
015 QUERY SESSION
```

Example: Update a specific line a script

Using the script from the prior example, change line 010 so that it processes the QUERY STGPPOOL command instead of the QUERY PROCESS command:

```
update script qsample "query stgpool" line=010
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
```

```
010 QUERY STGPOOL
015 QUERY SESSION
```

Example: Insert a command in the middle of a script

Using the script from the prior example, insert a new command line (QUERY NODE) after the QUERY STATUS command line in the QSAMPLE script:

```
update script qsample "query node"
line=007
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
007 QUERY NODE
010 QUERY STGPOOL
015 QUERY SESSION
```

Related commands

Table 1. Commands related to UPDATE SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the IBM Spectrum Protect server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.

Related tasks:

Running commands in parallel or serially
Including logic flow statements in a script
Performing tasks concurrently on multiple servers
Defining a server script

Related reference:

Return codes for use in IBM Spectrum Protect scripts

UPDATE SERVER (Update a server defined for server-to-server communications)

Use this command to update a server definition.

Restriction: If this server is a source server for a virtual volume operation, changing any of these values can affect the ability of the source server to access and manage the data that is stored on the corresponding target server. Changing the server name by using the SET SERVERNAME command might have additional implications, varying by operating system. The following are some examples:

- Passwords might be invalidated
- Device information might be affected
- Registry information about Windows operating systems might change

Privilege class

To issue this command, you must have system privilege.

Syntax for:

- **Enterprise configuration**

- Enterprise event logging
- Command routing
- Storage agent
- Node replication source and target servers
- **AIX** | **Linux** | **z/OS®** media server

```

>>-UPDate--SERver--server_name----->
>--+-----+----->
  '-SERVERPAssword---password-'
>--+-----+----->
  '-HLAddress---ip_address-'  '-LLAddress---tcp_port-'
>--+-----+----->
  '-COMMmethod---TCPIP-'  '-URL---url-'
>--+-----+----->
  '-ALLOWReplace---+Yes-+'
                               '-No--'
>--+-----+----->
  '-DESCRiption---description-'  '-FORCESync---+Yes-+'
                                       '-No--'
>--+-----+----->
  | (1) |
  '-VALIDateprotocol---+No-+-'
                               '-All-'
>--+-----+----->
  '-SSL---+No-+-'
                '-Yes-'

.-SESSIONSECurity---TRANSitional----.
>--+-----+----->
  '-SESSIONSECurity---+STRict-----+'
                               '-TRANSitional-'

.-TRANSFERMethod---Tcpi-----.
>--+-----+----->
  '-TRANSFERMethod---+Tcpi-----+'
                               | (2) |
                               '-Fasp-----'

```

Notes:

1. The VALIDATEPROTOCOL parameter is deprecated and applies only to storage agent definitions.
2. **Linux** The TRANSFERMETHOD parameter is available only on Linux x86_64 operating systems.

Syntax for virtual volumes

```

>>-UPDate--SERver--server_name----->
                               '-PAssword---password-'
>--+-----+----->
  '-HLAddress---ip_address-'  '-LLAddress---tcp_port-'
>--+-----+----->
  '-COMMmethod---TCPIP-'  '-URL---url-'
>--+-----+----->
  '-DELgraceperiod---days-'  '-NODEName---node_name-'

                               .-SESSIONSECurity---TRANSitional----.
>--+-----+----->
  '-SSL---Yes-'  '-SESSIONSECurity---+STRict-----+'
                                       '-TRANSitional-'

```

```
>-----+-----+-----+-----+-----+-----+-----+-----+-----+-----<
'-FORCESync---+---Yes---' '-DEScRiption---description-'
'-No--'
```

Parameters

server_name (Required)

Specifies the name of the server to be updated. This parameter is required.

PAssword

Specifies the password that is used to sign on to the target server for virtual volumes. This parameter is optional. If you specify a password, the minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

SERVERPAssword

Specifies the server password, which is used for enterprise configuration, command routing, and server-to-server event logging functions. The password must match the server password that is set by the SET SERVERPASSWORD command. This parameter is optional. The minimum length of the password is 8 characters unless a different value is specified by using the SET MINPWLENGTH command. The maximum length of the password is 64 characters.

HLAddress

Specifies the IP address (in dotted decimal format) of the server. This parameter is optional.

LLAddress

Specifies the low-level address of the server. This address is usually the same as the address in the TCPPOINT server option of the target server. When SSL=YES, the port must already be designated for SSL communications on the target server.

COMMmethod

Specifies the communication method that is used to connect to the server. This parameter is optional.

URL

Specifies the URL address that is used to access this server from the Administration Center. The parameter is optional.

DELgraceperiod

Specifies a number of days that an object remains on the target server after it was marked for deletion. You can specify a value 0 - 9999. The default is 5. This parameter is optional.

NODENAME

Specifies a node name to be used by the server to connect to the target server. This parameter is optional.

DEScRiption

Specifies a description of the server. This parameter is optional. The description can be up to 255 characters. Enclose the description in quotation marks if it contains blank characters. To remove an existing description, specify a null string (").

FORCESync

Specifies whether to reset the server verification key when the source server next signs on to the target server. A valid verification key enables a source server to put objects on the target server, manage the grace deletion period value, and update the password, if the current password is known and the verification key matches. The parameter is optional. You can specify one of the following values:

Yes

Specifies that a new verification key will be sent to and accepted by the target server if a valid password is received.

No

Specifies that a new verification key will not be sent to the target server.

VALIDateprotocol (deprecated)

Specifies whether a cyclic redundancy check validates the data sent between the storage agent and the IBM Spectrum Protect™ server. The parameter is optional. The default is NO.

Important: Beginning with IBM Spectrum Protect Version 8.1.2 and Tivoli® Storage Manager Version 7.1.8, validation that is enabled by this parameter is replaced by the TLS 1.2 protocol, which is enforced by the SESSIONSECURITY parameter. The VALIDATEPROTOCOL parameter is ignored. Update your configuration to use the SESSIONSECURITY parameter.

ALLOWReplace

Specifies whether a server definition that was defined by a managed server can be replaced with a definition from the configuration manager. This parameter is optional. You can specify one of the following values:

Yes

Specifies that a server definition can be replaced by a definition from the configuration manager.

No

Specifies that a server definition cannot be replaced by the definition from the configuration manager.

SSL

Specifies the communication mode of the server.

Important: Beginning with IBM Spectrum Protect V8.1.2 and Tivoli Storage Manager V7.1.8, SSL is used to encrypt some communication with the specified server even when you specify NO.

The following conditions and considerations apply when you specify the SSL parameter:

- Before starting the servers, self-signed certificates of the partner servers must be in the key database file (cert.kdb) of each of the servers.
- You can define multiple server names with different parameters for the same target server.

You can specify one of the following values:

No

Specifies an SSL session for all communication with the specified server, except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure.

Yes

Specifies an SSL session for all communication with the specified server, even when the server is sending and receiving object data.

SESSIONSECURITY

Specifies whether the server that you are defining must use the most secure settings to communicate with an IBM Spectrum Protect server. This parameter is optional.

You can specify one of the following values:

STRICT

Specifies that the strictest security settings are enforced for the server that you are defining. The STRICT value uses the most secure communication protocol available, which is currently TLS 1.2. The TLS 1.2 protocol is used for SSL sessions between the specified server and an IBM Spectrum Protect server.

To use the STRICT value, the following requirements must be met to ensure that the specified server can authenticate with the IBM Spectrum Protect server:

- Both the server that you are defining and the IBM Spectrum Protect server must be using IBM Spectrum Protect software that supports the SESSIONSECURITY parameter.
- The server that you are defining must be configured to use the TLS 1.2 protocol for SSL sessions between itself and the IBM Spectrum Protect server.

Servers set to STRICT that do not meet these requirements are unable to authenticate with the IBM Spectrum Protect server.

TRANSITIONAL

Specifies that the existing security settings are enforced for the server. This is the default value. This value is intended to be used temporarily while you update your security settings to meet the requirements for the STRICT value.

If SESSIONSECURITY=TRANSITIONAL and the server has never met the requirements for the STRICT value, the server will continue to authenticate by using the TRANSITIONAL value. However, after a server meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. Then, the server can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT. In addition, after a server successfully authenticates by using a more secure communication protocol, the server can no longer authenticate by using a less secure protocol. For example, if a server that is not using SSL is updated and successfully authenticates by using TLS 1.2, the server can no longer authenticate by using no SSL protocol or TLS 1.1. This restriction also applies when you use functions such as virtual volumes, command routing, or server-to-server export, when a node or administrator authenticates to the IBM Spectrum Protect server as a node or administrator from another server.

Linux TRANSFERMethod

Linux Specifies the method that is used for server-to-server data transfer. This parameter is optional. You can specify one of the following values:

Tcpip

Specifies that TCP/IP is used to transfer data. This is the default.

Fasp

Specifies that Aspera® Fast Adaptive Secure Protocol (FASP®) technology is used to transfer data. Aspera FASP technology can help you optimize data transfer in a wide area network (WAN).

Restrictions:

- Before you enable Aspera FASP technology, determine whether the technology is appropriate for your system environment and install the appropriate licenses. For instructions, see Determining whether Aspera FASP technology can optimize data transfer in your system environment. If the licenses are missing or expired, data transfer operations fail.
- If WAN performance meets your business needs, do not enable Aspera FASP technology.
- If you specify TRANSFERMETHOD=FASP on the PROTECT STGPOOL or REPLICATE NODE command, that value overrides the TRANSFERMETHOD parameter on the DEFINE SERVER and UPDATE SERVER commands.

Example: Update a deletion grace period for a server

Update the definition of SERVER2 to specify that objects remain on the target server for 10 days after they were marked for deletion.

```
update server server2 delgraceperiod=10
```

Example: Update the URL for a server

Update the definition of NEWSERVER to specify its URL address to be http://newserver:1580/.

```
update server newserver url=http://newserver:1580/
```

Example: Update all servers to communicate with an IBM Spectrum Protect server by using strict session security

Update the definition of all servers to use the strictest security settings to authenticate with the IBM Spectrum Protect server.

```
update server * sessionsecurity=strict
```

Related commands

Table 1. Commands related to UPDATE SERVER

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DEVCLASS	Deletes a device class.
DELETE FILESPACE	Deletes data associated with client file spaces. If a file space is part of a collocation group and you remove the file space from a node, the file space is removed from the collocation group.
DELETE SERVER	Deletes the definition of a server.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY SERVER	Displays information about servers.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE NODE	Changes the attributes that are associated with a client node.

UPDATE SERVERGROUP (Update a server group description)

Use this command to update the description of a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-UPDate SERVERGroup--group_name----->
>--DESCRiption---description-----<<
```

Parameters

group_name (Required)

Specifies the server group to update.

DESCRiption (Required)

Specifies a description of the server group. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

Example: Update the description of a server group

Update the description of the server group named WEST_COMPLEX to "Western Region Complex".

```
update servergroup west_complex
description="western region complex"
```

Related commands

Table 1. Commands related to UPDATE SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.

UPDATE SPACETRIGGER (Update the space triggers)

Use this command to update settings for triggers that determine when and how the server resolves space shortages in storage pools that use sequential-access FILE and random-access DISK device classes.

For storage pools with a parameter RECLAMATIONTYPE=SNAPLOCK, space triggers are not enabled.

Important: Space trigger functions and storage pool space calculations take into account the space remaining in each directory. Ideally, you associate each directory with a separate file system. If you specify multiple directories for a device class and the directories reside in the same file system, the server calculates space by adding values representing the space remaining in each directory. These space calculations will be inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the wrong storage pool and run out of space prematurely. For space triggers, an inaccurate calculation might result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled. If a trigger is disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by specifying the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

See the DEFINE SPACETRIGGER command for more information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
>>-UPDate SPACETrigger--STG--+-----+----->
                                     '-Fullpct--==--percent-'
>--+-----+----->
   '-SPACEexpansion--==--percent-'
>--+-----+----->
   '-EXPansionprefix--==--prefix-'
>--+-----+----->>
   '-STGPOOL--==--storage_pool_name-'
```

Parameters

STG (Required)

Specifies a storage pool space trigger

Fullpct

This parameter specifies the utilization percentage of the storage pool.

When this value is exceeded, the space trigger creates new volumes.

You can determine storage pool utilization by issuing the QUERY STGPOOL command with FORMAT=DETAILED. The percentage of storage pool utilization for the storage pool is displayed in the field "Space Trigger Util." The calculation for this percentage does not include potential scratch volumes. The calculation for the percentage utilization used for migration and reclamation, however, does include potential scratch volumes.

SPACEexpansion

For space triggers for sequential-access FILE-type storage pools, this parameter is used in determining the number of additional volumes that are created in the storage pool. Volumes are created using the MAXCAPACITY value from the storage pool's device class. For space triggers for random-access DISK storage pools, the space trigger creates a single volume using the EXPANSIONPREFIX.

EXPansionprefix

This specifies the prefix that the server uses to create new storage pool files. This parameter is optional and applies only to random-access DISK device classes. The default prefix is the server installation path.

The prefix can include one or more directory separator characters, for example:

AIX | **Linux**

```
/opt/tivoli/tsm/server/bin/
```

Windows

```
c:\program files\tivoli\tsm\
```

AIX | **Linux**

You can specify up to 250 characters. If you specify a prefix that is not valid, automatic expansion can fail.

Windows

You can specify up to 200 characters. If the server is running as a Windows service, the default prefix is the c:\wnnt\system32 directory. If you specify a prefix that is not valid, automatic expansion can fail.

This parameter is not valid for space triggers for sequential-access FILE storage pools. Prefixes are obtained from the directories specified with the associated device class.

STGPOOL

Specifies the storage pool associated with this space trigger. If the STGPOOL parameter is not specified, the default storage pool space trigger is updated.

This parameter does not apply to storage pools with the parameter RECLAMATIONTYPE=SNAPLOCK.

Example: Increase the amount of space for a storage pool

Increase the amount of space in a storage pool by 50 percent when it is filled to 80 percent utilization of existing volumes. Space will be created in the directories associated with the device class.

DBFREESPACE
Specifies the free space available in the database in gigabytes.

DBUSEDSPACE
Specifies the amount of database space that is used, in gigabytes.

ARCHIVELOGFREESPACE
Specifies the free space that is available in the archive log, in gigabytes.

STGPOOLUTIL
Specifies the storage pool utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

STGPOOLCAPACITY
Specifies the storage pool capacity in gigabytes.

AVGSTGPOOLUTIL
Specifies the average storage pool utilization percentage across all storage pools. The default warning threshold value is 80%, and the default error threshold value is 90%.

TOTSTGPOOLCAPACITY
Specifies the total storage pool capacity in gigabytes for all available storage pools.

TOTSTGPOOLS
Specifies the number of defined storage pools.

TOTRWSTGPOOLS
Specifies the number of defined storage pools that are readable or writeable.

TOTNOTRWSTGPOOLS
Specifies the number of defined storage pools that are not readable or writeable.

STGPOOLINUSEANDDEFINED
Specifies the total number of defined volumes that are in use.

ACTIVELOGUTIL
Specifies the current percent utilization of the active log. The default warning threshold value is 80%, and the default error threshold value is 90%.

ARCHLOGUTIL
Specifies the current utilization of the archive log. The default warning threshold value is 80%, and the default error threshold value is 90%.

CPYSTGPOOLUTIL
Specifies the percent utilization for a copy storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

PMRYSTGPOOLUTIL
Specifies the percent utilization for a primary storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

DEVCLASSPCTDRVOFFLINE
Specifies the percent utilization of drives that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDRVPOLLING
Specifies the drives polling, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTLIBPATHSOFFLINE
Specifies the library paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTPATHSOFFLINE
Specifies the percentage of device class paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDISKSNOTRW
Specifies the percentage of disks that are not writable for the disk device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDISKSUNAVAILABLE
Specifies the percentage of the disk volumes that are unavailable, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

FILEDEVCLASSPCTSCRUNALLOCATABLE
Specifies the percentage of scratch volumes that the server cannot allocate for a given non-shared file device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

Condition

Specify this value to change the condition of an existing threshold. This parameter is optional. Specify one of the following values:

- EXists
Creates a status monitoring indicator if the activity exists.
- GT
Creates a status monitoring indicator if the activity outcome is greater than the specified value.
- GE
Creates a status monitoring indicator if the activity outcome is greater than or equal to the specified value.
- LT
Creates a status monitoring indicator if the activity outcome is less than the specified value.
- LE
Creates a status monitoring indicator if the activity outcome is less than or equal to the specified value.
- EQual
Creates a status monitoring indicator if the activity outcome is equal to the specified value.

Value

Specify this parameter to change the value that is compared with the activity output for the specified condition. You can specify an integer in the range 0 - 999999999999999.

Status

Specify this value to change the status of the indicator that is created in status monitoring if the condition that is being evaluated passes. This parameter is optional. Specify one of the following values:

- Normal
Specifies that the status indicator has a normal status value.
- Warning
Specifies that the status indicator has a warning status value.
- Error
Specifies that the status indicator has an error status value.

Update an existing status threshold

Update a status threshold for average storage pool utility percentage by issuing the following command:

```
update statusthreshold avgstgpl "AVGSTGPOOLUTIL" value=90 condition=gt status=error
```

Related commands

Table 1. Commands related to UPDATE STATUSTHRESHOLD

Command	Description
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	Deletes a status monitoring threshold.
QUERY MONITORSTATUS (Query the monitoring status)	Displays information about monitoring alerts and server status settings.
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	Displays information about monitoring alerts and server status settings.
QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	Displays information about a status monitoring thresholds.
SET STATUSMONITOR (Specifies whether to enable status monitoring)	Specifies whether to enable status monitoring.
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	Specifies whether to enable client at-risk activity interval evaluation
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	Specifies the refresh interval for status monitoring.
SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	Specifies whether to use client at-risk skipped files as failure evaluation
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	Changes the attributes of an existing status monitoring threshold.

UPDATE STGPOOL (Update a storage pool)

Use this command to change a storage pool.

Restriction: If a client is using the simultaneous-write function and data deduplication, the data deduplication feature is disabled during backups to a storage pool.

The UPDATE STGPOOL command takes seven forms. The syntax and parameters for each form are defined separately.

Table 1. Commands related to UPDATE STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
COPY ACTIVATEDATA	Copies active backup data.
DEFINE COLLOGGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node or file space to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node or file space from a collocation group.
DELETE STGPOOL	Deletes a storage pool from server storage.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
MOVE MEDIA	Moves storage pool volumes that are managed by an automated library.
QUERY COLLOGGROUP	Displays information about collocation groups.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
SET DRMDBBACKUPEXPIREDDAYS	Specifies criteria for database backup series expiration.
SHRED DATA	Manually starts the process of shredding deleted data.
UPDATE COLLOGGROUP	Updates the description of a collocation group.

- UPDATE STGPOOL (Update a cloud-container storage pool)
Use this command to update a container storage pool in a cloud environment. Cloud storage pools are not supported on Linux on System z®.
- UPDATE STGPOOL (Update a directory-container storage pool)
Use this command to update a directory-container storage pool.
- UPDATE STGPOOL (Update a container-copy storage pool)
Use this command to update a container-copy storage pool.
- UPDATE STGPOOL (Update a primary random access storage pool)
Use this command to update a random access storage pool.
- UPDATE STGPOOL (Update a primary sequential access pool)
Use this command to update a primary sequential access storage pool.
- UPDATE STGPOOL (Update a copy sequential access storage pool)
Use this command to update a copy sequential access storage pool.
- UPDATE STGPOOL (Update an active-data sequential access)
Use this command to update an active-data pool.

UPDATE STGPOOL (Update a cloud-container storage pool)

Use this command to update a container storage pool in a cloud environment. Cloud storage pools are not supported on Linux on System z®.

The preferred way to define and configure a cloud-container storage pool is to use the Operations Center. For instructions and tips for the Operations Center and the command-line interface, see [Configuring a cloud-container storage pool for data storage](#).

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```
>>-UPDate STGpool--pool_name--+-+-----+----->
                                '-DEscription--==--description-'
>--+-----+----->
|                                     (1) |
'-CLOUDType--==--+Swift-----+-----'
      +-IBMCloudswift-+
      '-VlSwift-----'
>--+-----+----->
'-CLOUDUrl--==--cloud_url-'
>--+-----+----->
|                                     (2) |
'-IDentity--==--cloud_identity-----'
>--+-----+----->
'-PAssword--==--password-'
>--+-----+----->
'-CLOUDLocation--==--+OFFpremise-+-'
                        '-ONpremise--'
>--+-----+----->
|                                     (3) |
'-BUCKETName--==--bucket_name-----'
>--+-----+----->
'-ACCess--==--+READWrite---+-'
      +-READOnly-----+
      +-UNAVailable-+
      '-DESTroyed---'
>--+-----+----->
'-MAXWriters--==--+NOLimit-----+-'
                        '-maximum_writers-'
>--+-----+----->
'-REUsedelay--==--days-'
>--+-----+-----><
|                                     .-COMPReSSion--==--Yes-----. |
'-ENCRypt--==--+Yes-+-+-----+-----'
      '-No--'   '-COMPReSSion--==--+Yes-+-'
                        '-No--'
```

Notes:

1. CLOUDTYPE=S3 and CLOUDTYPE=AZURE cannot be changed.
2. For Azure storage pools, it is not necessary to specify the IDENTITY parameter.
3. This parameter is valid only if you specify CLOUDTYPE=S3.

Parameters

pool_name (Required)

Specifies the storage pool to update. This parameter is required.

DEscription

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters. To remove an existing description, specify a null string ("").

CLOUDType

Specifies the type of cloud environment where you are configuring a storage pool. This parameter is optional. Specify one of the following values:

IBMCloudswift

Specifies that the storage pool uses an IBM® Cloud cloud computing system with an OpenStack Swift cloud computing system.

SWift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 2 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol it is using.

V1Swift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 1 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol it is using.

Restriction: If you used the DEFINE STGPOOL command to define a storage pool with CLOUDTYPE=S3 (Simple Storage Service) or CLOUDTYPE=AZURE, you cannot change to a different cloud type by using the UPDATE STGPOOL command. Additionally, you cannot change the following cloud types by using the UPDATE STGPOOL command:

- A non-S3 storage pool to S3
- A non-Azure storage pool to Azure

CLOUDUrl

Specifies the URL of the cloud environment where you are configuring the storage pool. Based on your cloud provider, you can use a region endpoint URL, an accesser IP address, a public authentication endpoint, or a similar value for this parameter. Be sure to include the protocol, such as `https://` or `http://`, at the beginning of the URL. The maximum length of the web address is 870 characters. The CLOUDURL parameter is not validated until the first backup begins. For more information about how to locate these values, select your cloud service provider from the list on the Configuring a cloud-container storage pool for data storage page.

Tip: To use more than one IBM Cloud Object Storage accesser, list the accesser IP addresses separated by a vertical bar (|), with no spaces, such as in the following example:

```
CLOUDURL=<accesser_URL1>|<accesser_URL2>|<accesser_URL3>
```

Use multiple accessers to improve performance. If you are using the IBM Cloud S3 solution, only one accesser is needed.

Identity

Specifies the user ID for the cloud that is specified in the STGTYPE=CLOUD parameter. This parameter is required for all supported cloud computing systems except Azure. If you specified CLOUDTYPE=AZURE, do not specify the IDENTITY parameter. Based on your cloud provider, you can use an access key ID, a user name, a tenant name and user name, or a similar value for this parameter. The maximum length of the user ID is 255 characters.

PAssword (Required)

Specifies the password for the cloud that is specified in the STGTYPE=CLOUD parameter. Based on your cloud provider, you can use a shared access signature (SAS) token, secret access key, an API key, a password, or a similar value for this parameter. This parameter is required. The maximum length of the password is 255 characters. The IDENTITY and PASSWORD parameters are not validated until the first backup begins.

CLOUDLocation

Specifies the physical location of the cloud that is specified in the CLOUD parameter. This parameter is optional. You can specify one of the following values:

- Offpremise
- ONpremise

BUCKETName

Specifies the name for an Amazon Web Services (AWS) bucket or IBM Cloud Object Storage vault to use with this storage pool. AWS buckets and IBM Cloud Object Storage vaults are used in the same manner as containers in a cloud-container storage pool. This parameter is optional, and is valid only if this storage pool has a cloud type of S3. If the name that you specify does not exist, the server creates a bucket or vault with the specified name before using the bucket or vault. Follow

the naming restrictions for your cloud provider when specifying this parameter. Review the permissions for the bucket or vault and ensure that the credentials for this storage pool have permission to read, write, list, and delete objects in this bucket or vault.

Restriction: You cannot change the bucket or vault if any cloud containers exist in this storage pool.

ACCess

Specifies how client nodes and server processes access the storage pool. This parameter is optional. You can specify one of the following values:

READWrite

Specifies that client nodes and server processes can read and write to the storage pool.

READOnly

Specifies that client nodes and server processes can read only from the storage pool.

UNAVailable

Specifies that client nodes and server processes cannot access the storage pool. As a result, backups and restore fail for this storage pool. You can use this value to specify that the cloud service provider is temporarily unavailable.

DESTroyed

Specifies that client nodes and server processes cannot access the storage pool because the cloud service provider is permanently unavailable. Backups and restores fail for this storage pool, but any attempts to delete objects and containers from this storage pool finish successfully.

MAXWriters

Specifies the maximum number of writing sessions that can run concurrently on the storage pool. Specify a maximum number of writing sessions to control the performance of the cloud storage pool from negatively impacting other system resources. This parameter is optional. You can specify one of the following values:

NOLimit

Specifies that no maximum size limit exists for the number of writers that you can use. This value is the default.

maximum_writers

Limits the maximum number of writers that you can use. Specify an integer in the range 1 - 99999.

REUsedelay

Specifies the number of days that must elapse after all deduplicated extents are removed from a cloud storage pool. This parameter controls the duration that deduplicated extents are associated with a cloud storage pool. When the value that is specified for the parameter expires, the deduplicated extents are deleted from the cloud storage pool. This parameter is optional. You can specify one of the following values:

1

Specifies that deduplicated extents are deleted from a cloud storage pool after one day.

days

You can specify an integer in the range 0 - 9999.

Tip: Set this parameter to a value that is greater than the number specified for the SET DRMDBBACKUPEXPIREDAYS command. By setting this parameter to a higher value, you can ensure that when you restore the database to an earlier level, the references to files in the storage pool are still valid.

ENCRypt

Specifies whether the server encrypts client data before it writes it to the storage pool. You can specify the following values:

Yes

Specifies that client data is encrypted by the server.

No

Specifies that client data is not encrypted by the server.

This parameter is optional. The default depends on the physical location of the cloud, which is specified by the CLOUDLOCATION parameter. If the cloud is off premise, the server encrypts data by default. If the cloud is on premises, the server does not encrypt data by default.

COMPRession

Specifies whether data is compressed in the storage pool. This parameter is optional. You can specify one of the following values:

No

Specifies that data is not compressed in the storage pool.

Yes

Specifies that data is compressed in the storage pool. This is the default.

Example 1: Update a cloud storage pool to specify a maximum number of data sessions

Update a cloud storage pool that is named STGPOOL1 and specify a maximum of 10 data sessions.

```
update stgpool stgpool1 maxwriters=10
```

Example 2: Update the description of a cloud-container storage pool

Update a cloud-container storage pool that is named STGPOOL2. Remove the existing description from the storage pool.

```
update stgpool stgpool2 cloudurl=http://123.234.123.234:5000/v2.0
identity=admin:admin password=protect8991 description=""
```

Related tasks:

Configuring a cloud-container storage pool for data storage

AIX Linux Windows

UPDATE STGPOOL (Update a directory-container storage pool)

Use this command to update a directory-container storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```
>>-Update STGpool--pool_name--+-----+----->
                                     '-DEScRiption--description-'
                                     +-----+----->
.-ACCess---READWrite-----
>--+-----+----->
  '-ACCess---+READWrite---+'
                    +READOnly---+
                    '-UNAVailable-'
                                     +-----+----->
.-MAXSIZe---NOLimit-----
>--+-----+----->
  '-MAXSIZe---+maximum_file_size+-'
                    '-NOLimit-----'
                                     +-----+----->
.-MAXWriters---NOLimit-----
>--+-----+----->
  '-MAXWriters---+maximum_writers+-'
                    '-NOLimit-----'
                                     +-----+----->
  '-NEXTstgpool---pool_name-'
                                     +-----+----->
  '-PROTECTstgpool---target_stgpool-'
                                     +-----+----->
|                                     .-,-----|
|                                     V           ||
  '-PROTECTLOCalstgpoools---local_target_stgpool+-'
                                     +-----+----->
.-REUsedelay---1-----
>--+-----+----->
  '-REUsedelay---days-'  '-ENCRypt---+Yes+-'
                                     '-No--'
                                     +-----+----->
.-COMPRession---Yes-----
>--+-----+----->>
  '-COMPRession---+Yes+-'
```

Parameters

pool_name (Required)

Specifies the storage pool to update. This parameter is required. The maximum length of the name is 30 characters.

DESCRiption

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how client nodes and server processes access files in the storage pool. This parameter is optional. You can specify one of the following values:

READWrite

Specifies that client nodes and server processes can read and write to the storage pool. This is the default.

READOnly

Specifies that client nodes and server processes can only read from the storage pool.

UNAVailable

Specifies that client nodes and server processes cannot access the storage pool.

MAXSIze

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is NOLIMIT. Specify one of the following values:

NOLimit

Specifies that there is no maximum size limit for physical files that are stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer in the range 1 - 999999, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 GB. Use one of the following scale factors:

Table 1. Scale factor
for the maximum file
size

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

Tip: If you do not specify a unit of measurement for the maximum file size, the value is specified in bytes.

When the physical size of the storage pool exceeds the MAXSIZE parameter, the following table shows where files are typically stored.

Table 2. The location of a file according to the file size and the pool that is specified

Pool that is specified	Result
No pool is specified as the next storage pool in the hierarchy.	The server does not store the file.
A pool is specified as the next storage pool in the hierarchy.	The server stores the file in the storage pool that you specified.

Tip: If you also specify the NEXTstgpool parameter, update one storage pool in your hierarchy to have no limit on the maximum file size by specifying the MAXSIZE=NOLimit parameter. When you have at least one pool with no size limit, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent during data deduplication processing, the server considers the size of the data deduplication process to be the file size. If the total size of all files in the process is larger than the maximum size limit, the server does not store the files in the storage pool.

MAXWriters

Specifies the maximum number of I/O threads that can run concurrently on the storage pool. Specify a maximum number of I/O threads to control the number of I/O threads that are written simultaneously to the directory-container storage pool. This parameter is optional. As a best practice, use the default value of NOLIMIT. You can specify one of the following values:

NOLimit

Specifies that no maximum number of I/O threads are written to the storage pool.

maximum_writers

Limits the maximum number of I/O threads that you can use. Specify an integer in the range 1 - 99999.

NEXTstgpool

Specifies the name of a random-access or primary sequential storage pool to which files are stored when the directory-container storage pool is full. This parameter is optional.

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.
- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.

PROTECTstgpool

Specifies the name of the directory-container storage pool on the target server where the data is backed up when you use the PROTECT STGPOOL command for this storage pool. This parameter is optional.

PROTECTLOCALstgpools

Specifies the name of the container-copy storage pool on a local device where the data is backed up. This container-copy storage pool will be a local target storage pool when you use the PROTECT STGPOOL command. You can specify a maximum of two container-copy storage pool names to update. Separate multiple names with commas and no intervening spaces. The maximum length of each name is 30 characters. This parameter is optional.

To add or remove container-copy storage pools, specify the container-copy storage pool names to include. For example, if the existing container-copy storage pool includes COPY1 and you want to add COPY2, specify PROTECTLOCALSTGPOOLS=COPY1,COPY2. To remove all existing container-copy storage pools that are associated with the primary storage pool, specify a null string (""). For example, COPYSTGPOOLS="".

REUsedelay

Specifies the number of days that must elapse before all deduplicated extents are removed from a directory-container storage pool. This parameter controls the duration that deduplicated extents are associated with a directory-container storage pool. When the value that is specified for the parameter expires, the deduplicated extents are deleted from the directory-container storage pool. The default is 1. Specify one of the following values:

days

Specify an integer in the range 0 - 9999.

1

Specifies that deduplicated extents are deleted from a directory-container storage pool after one day.

Tip: Set this parameter to a value greater than the number that is specified as your database backup period to ensure that data extents are still valid when you restore the database to another level.

ENCRypt

Specifies whether the server encrypts client data before the server writes the data to the storage pool. You can specify the following values:

Yes

Specifies that client data is encrypted by the server.

No

Specifies that client data is not encrypted by the server.

COMPReSSion

Specifies whether data is compressed in the storage pool. This parameter is optional. You can specify one of the following values:

No

Specifies that data is not compressed in the storage pool.

Yes

Specifies that data is compressed in the storage pool. This is the default.

Example: Update a storage pool to specify a maximum number of data sessions

Update a storage pool that is named STGPOOL1 and specify a maximum of 10 data sessions.

```
update stgpool stgpool1 maxwriters=10
```

Example: Update a storage pool to specify the maximum size

Update a storage pool that is named STGPOOL2. The storage pool specifies the maximum file size that the server can store in the storage pool as 100 megabytes.



```
update stgpool stgpool2 maxsize=100M
```

Example: Update the description of a storage pool

Update a storage pool that is named STGPOOL3. Remove the existing description from the storage pool.

```
update stgpool stgpool3 description=""
```

Table 3. Commands related to UPDATE STGPOOL

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
QUERY CONTAINER	Displays information about a container.
QUERY STGPOOL	Displays information about storage pools.
REPAIR STGPOOL	Repairs a directory-container storage pool.
 UPDATE STGPOOLDIRECTORY	Changes the attributes of a storage pool directory.
	

UPDATE STGPOOL (Update a container-copy storage pool)

Use this command to update a container-copy storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax

```
>>-UPDate STGpool--pool_name--+-----+----->
                                     '-MAXSCRatch---number-'

>--+-----+----->
   '-DESCRiption---description-'

>--+-----+----->
   '-ACCess---+READWrite---+'
                   +-READOnly---+
                   '-UNAVailable-'

>--+-----+--+-----+----->
   '-PROTECTPRocess---number-' '-REClaim---percent-'

>--+-----+----->
```

```
'-RECLAIMLimit-----+--NOLimit---+-'
      '-vol_limit-'
>---+-----+-----><
      '-REUsedelay----days-'
```

Parameters

pool_name (Required)

Specifies the name of the storage pool to be updated.

MAXSCRatch

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer in the range 0 - 100000000. If the server can request scratch volumes as needed, you do not have to define each volume to be used.

The value of this parameter is used to estimate the total number of volumes that are available in the storage pool and the corresponding estimated capacity for the storage pool.

DESCription

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how server processes such as storage-pool protection and repair can access data in the storage pool. This parameter is optional. You can specify one of the following values:

READWrite

Specifies that the server can read and write to volumes in the storage pool.

READOnly

Specifies that the server can only read volumes in the storage pool. The server can use data in the storage pool to restore extents to directory-container storage pools. No operations that write to the container-copy storage pool are allowed.

UNAVailable

Specifies that the server cannot access data that is stored on volumes in the storage pool.

PROTECTPProcess

Specifies the maximum number of parallel processes that are used when you issue the PROTECT STGPOOL command to copy data to this pool from a directory-container storage pool. This parameter is optional. Enter a value in the range 1 - 20.

The time that is required to complete the copy operation might be decreased by using multiple, parallel processes. However, in some cases when multiple processes are running, one or more of the processes must wait to use a volume that is already in use by a different process.

When you select this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a tape volume, the server uses a mount point and a drive. The number of available mount points and drives depends on the mount limit of the device class for the storage pool, and on other server and system activity.

If you use the preview option on the PROTECT STGPOOL command, only one process is used and no mount points or drives are needed.

REClaim

Specifies when a volume becomes eligible for reclamation and reuse. Specify eligibility as the percentage of a volume's space that is occupied by extents that are no longer stored in the associated directory-container storage pool. Reclamation moves any extents that are still stored in the associated directory-container storage pool from eligible volumes to other volumes. Reclamation occurs only when a PROTECT STGPOOL command stores data into this storage pool.

This parameter is optional. You can specify an integer in the range 1 - 100. The value 100 specifies that volumes in this storage pool are not reclaimed.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

By setting the reclaim value to 50 percent or greater, data that is moved from two reclaimed volumes uses no more than the equivalent of one new volume.

Use caution when you use reclamation with container-copy storage pools that have offsite volumes. When an offsite volume becomes eligible for reclamation, in effect the server moves the extents on the volume back to the onsite location. If a disaster occurs onsite, the server can obtain extents from the offsite volume if the restored database refers to extents on the offsite volume. Therefore, for disaster recovery purposes, ensure that you schedule database backups to run after storage pool protection schedules and DRM move schedules have run, and ensure that all database backup volumes are taken offsite along with the DRM volumes.

Tip: Set different reclamation values for offsite container-copy storage pools and onsite container-copy storage pools. Because container-copy storage pools store deduplicated data, the data extents are spread across multiple tape volumes. When you choose a reclamation threshold for an offsite copy, carefully consider the number of available mount points and the number of tape volumes that you must retrieve if a disaster occurs. Setting a higher threshold means that you must retrieve more volumes than you would if your reclamation value was lower. Using a lower threshold reduces the number of mount points that are required in a disaster. The preferred method is to set the reclamation value for offsite copies to 60, and for onsite copies, in the range 90 - 100.

RECLAIMLimit

Specifies the maximum number of volumes that the server reclaims when you issue the PROTECT STGPOOL command and specify the RECLAIM=YESLIMITED or RECLAIM=ONLYLIMITED option. This parameter is valid only for container-copy storage pools. This parameter is optional. You can specify one of the following values:

NOLimit

Specifies that all volumes in the container-copy storage pool are processed for reclamation.

vol_limit

Specifies the maximum number of volumes in the container-copy storage pool that are reclaimed. The value that you specify determines how many new scratch tapes are available after reclamation processing completes. You can specify a number in the range 1 - 100000.

REUsedelay

Specifies the number of days that must elapse after all extents are deleted from a volume before the volume can be rewritten or returned to scratch status. This parameter is optional. You can specify an integer in the range 0 - 9999. A value of 0 means that a volume can be rewritten or returned to scratch status as soon as all the extents are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to extents in the storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. If you use disaster recovery manager, the number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

Example: Update a container-copy storage pool to delay volume reuse for 30 days

Update the storage pool that is named CONTAINER1_COPY2 to change the delay for volume reuse to 30 days.

```
update stgpool container1_copy2 reusedelay=30
```

Example: Update a container-copy storage pool to limit the number of reclaimed tape volumes to 10

Update the storage pool that is named CONTAINER1_COPY2 to change the reclaim limit to 10 volumes.

```
update stgpool container1_copy2 reclaimlimit=10
```

Table 1. Commands related to UPDATE STGPOOL (Update a container-copy storage pool)

Command	Description
DEFINE STGPOOL (container-copy)	Define a container-copy storage pool that stores copies of data from a directory-container storage pool.
PROTECT STGPOOL	Protects a directory-container storage pool.
QUERY STGPOOL	Displays information about storage pools.
REPAIR STGPOOL	Repairs a directory-container storage pool.
UPDATE STGPOOL (directory-container)	Update a directory-container storage pool.

UPDATE STGPOOL (Update a primary random access storage pool)

Use this command to update a random access storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax

```
>>-UPDate STGpool--pool_name--+-+-----+-----+----->
                                     '-DESCRiption--==--description-'
>--+-----+-----+----->
   '-ACCess--==--READWrite--+-+'
       +-READOnly--+-+
       '-UNAVailable-'
>--+-----+-----+----->
   '-MAXSize--==--maximum_file_size--+'
       '-NOLimit-----'
>--+-----+-----+----->
   '-CRCData--==--Yes--+'   '-NEXtstgpool--==--pool_name-'
       '-No--'
>--+-----+-----+----->
   '-HIghmig--==--percent-'   '-LOWmig--==--percent-'
>--+-----+-----+----->
   '-CACHe--==--Yes--+'   '-MIGPRocess--==--number-'
       '-No--'
>--+-----+-----+----->
   '-MIGDelay--==--days-'   '-MIGContinue--==--Yes--+'
                                   '-No--'
>--+-----+-----+----->
   '-AUTOCopy--==--None-----+'
       +-CLient--+-+
       +-MIGRation--+-+
       '-All-----'
>--+-----+-----+----->
   |                                     .-,----- . |
   |                                     v          | |
   '-COPYSTGpools--==--coppoolname--+'
>--+-----+-----+----->
   '-COPYContinue--==--Yes--+'
       '-No--'
>--+-----+-----+----->
   |                                     .-,----- . |
   |                                     v          | |
   '-ACTIVEDATApools--==--active-data_pool_name--+'
   .-SHRED--==--0-----
>--+-----+-----+-----><
   '-SHRED--==--overwrite_count-'
```

Parameters

pool_name (Required)

Specifies the storage pool to update. This parameter is required.

DESCRiption

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. You can specify the following values:

READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. You can specify the following values:

NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer from 1 to 999999 terabytes, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

The client estimates the size of files that are sent to the server. The client estimate is used rather than the actual amount of data that is sent to the server. Client options, such as deduplication, compression, and encryption, can cause the actual amount of data that is sent to the server to be larger or smaller than the size estimate. For example, the compression of a file might be smaller in size than the estimate, thus sending less data than the estimate. Furthermore, a binary file might be larger in size after the compression processing, thus sending more data than the estimate.

See the following table for information about where a file is stored when its size exceeds the MAXSIZE parameter.

Table 1. Where a file is stored according to the file size and the pool that is specified

File size	Pool specified	Result
Exceeds the maximum size	No pool is specified as the next storage pool in the hierarchy	The server does not store the file
	A pool is specified as the next storage pool in the hierarchy	The server stores the file in the next storage pool that can accept the file size

If you specify the next storage pool parameter, define one storage pool in your hierarchy to have no limit on the maximum file size. By having no limit on the size for at least one pool, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent in a single transaction, the server considers the size of the transaction to be the file size. If the total size of all files in the transaction is larger than the maximum size limit, the server does not store the files in the

storage pool.

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more expenditure is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

NEXTstgpool

Specifies a primary storage pool to which files are migrated. This parameter is optional.

To remove an existing storage pool from the storage hierarchy, specify a null string ("") for this value.

If you do not specify a next storage pool, the following actions occur:

- The server cannot migrate files from this storage pool
- The server cannot store files that exceed the maximum size for this storage pool in another storage pool

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.
- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.

HIghmig

Specifies that the server starts migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer 0 - 100.

When the storage pool exceeds the high migration threshold, the server can start migration of files by node to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool.

LOWmig

Specifies that the server stops migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. You can specify an integer 0 - 99 for this optional parameter.

When migration is by node or file space, depending upon collocation, the level of the storage pool can fall below the value that you specified for this parameter. To empty the storage pool, set LOWMIG=0.

CAChe

Specifies whether the migration process leaves a cached copy of a file in this storage pool after you migrate the file to the next storage pool. This parameter is optional. You can specify the following values:

Yes

Specifies that caching is enabled.

No

Specifies that caching is disabled.

Using cache might improve your ability to retrieve files, but might affect the performance of other processes.

MIGPRocess

Specifies the number of processes that are used for migrating files from this storage pool. This parameter is optional. You can specify an integer 1 - 999.

During migration, these processes are run in parallel to provide the potential for improved migration rates.

Tips:

- The number of migration processes is dependent upon the following settings:
 - The setting of the MIGPROCESS parameter
 - The collocation setting of the next pool
 - The number of nodes or the number of collocation groups with data in the storage pool that is being migratedFor this example, `MIGPROCESS =6`, the next pool `COLLOCATE` parameter is `NODE`, but there are only two nodes with data on the storage pool. Migration processing consists of only two processes, not six. If the `COLLOCATE` parameter is `GROUP` group and both nodes are in the same group, migration processing consists of only one process. If the `COLLOCATE` parameter is `NO` or `FILESPACE` group, and each node has two file spaces with backup data, then migration processing consists of only four processes.
- When you specify this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool that is defined to the target storage pool.

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. To calculate a value to compare to the specified `MIGDELAY` value, the server counts the following items:

- The number of days that the file was in the storage pool
- The number of days, if any, since the file was retrieved by a client

The lesser of the two values are compared to the specified `MIGDELAY` value. For example, if all the following conditions are true, a file is not migrated:

- A file was in a storage pool for five days.
- The file was accessed by a client within the past three days.
- The value that is specified for the `MIGDELAY` parameter is four days.

This parameter is optional. You can specify an integer 0 - 9999. The default is 0, which means that you do not want to delay migration.

If you want the server to count the number of days that are based on when a file was stored and not when it was retrieved, use the `NORETRIEVEDATE` server option.

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

You can specify one of the following values:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

No

Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

AUTOCopy

Specifies when IBM Spectrum Protect™ runs simultaneous-write operations to copy storage pools and active-data pools. This parameter affects the following operations:

- Client store sessions
- Server import processes

- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the `COPYSTGPOLLS` parameter. Active-data pools are specified using the `ACTIVEDATAPOOLS` parameter.

You can specify one of the following values:

None

Specifies that the simultaneous-write function is disabled.

CLient

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

MIGRation

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

All

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

COPYSTGpools

Specifies the names of copy storage pools where the server simultaneously writes data. You can specify a maximum of three copy pool names that are separated by commas. Spaces between the names of the copy pools are not allowed. To add or remove one or more copy storage pools, specify the pool name or names that you want to include in the updated list. For example, if the existing copy pool list includes `COPY1` and `COPY2` and you want to add `COPY3`, specify `COPYSTGPOLLS=COPY1,COPY2,COPY3`. To remove all existing copy storage pools that are associated with the primary storage pool, specify a null string ("") for the value (for example, `COPYSTGPOLLS=""`).

When you specify a value for the `COPYSTGPOLLS` parameter, you can also specify a value for the `COPYCONTINUE` parameter. For more information, see the `COPYCONTINUE` parameter.

The combined total number of storage pools that are specified in the `COPYSTGPOLLS` and `ACTIVEDATAPOOLS` parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the `COPYCONTINUE` value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to copy storage pools for the following operations:

- Back up and archive operations by IBM Spectrum Protect backup-archive clients or application clients that are using the IBM Spectrum Protect API
- Migration operations by IBM Spectrum Protect for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

Restrictions: The simultaneous-write function is not supported for the following store operations:

- When the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is accepted.
- NAS backup operations. If the primary storage pool specified in the `DESTINATION` or `TOCDESTINATION` in the copy group of the management class has copy storage pools that are defined:
 - The copy storage pools are ignored
 - The data is stored into the primary storage pool only

Attention: The function that is provided by the `COPYSTGPOLLS` parameter is not intended to replace the `BACKUP STGPOL` command. If you use the `COPYSTGPOLLS` parameter, continue to use the `BACKUP STGPOL` command to ensure that the

copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

COPYContinue

Specifies how the server reacts to a copy storage pool write failure for any of the copy storage pools that are listed in the COPYSTGPOOLS parameter. This parameter is optional. When you specify the COPYCONTINUE parameter, either a COPYSTGPOOLS list must exist or the COPYSTGPOOLS parameter must also be specified.

You can specify the following values:

Yes

If the COPYCONTINUE parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

No

If the COPYCONTINUE parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the COPYCONTINUE parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The ACTIVEDATAPOOLS parameter is optional. Spaces between the names of the active-data pools are not allowed.

The combined total number of storage pools that are specified in the COPYSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool that is specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to active-data pools only during backup operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API.

Restrictions:

1. This parameter is available only to primary storage pools that use "NATIVE" or "NONBLOCK" data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Writing data simultaneously to active-data pools is not supported when you use LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is followed.
3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the TOCDESTINATION in the copy group of the management class has active-data pools that are defined:
 - The active-data pools are ignored
 - The data is stored into the primary storage pool only
4. You cannot use the simultaneous-write function with CENTERA storage devices.
5. Data that is being imported is not stored in active-data pools. After an import operation, use the COPY ACTIVEDATA command to store the imported data in an active-data pool.

Attention: The function that is provided by the `ACTIVEDATAPOOLS` parameter is not intended to replace the `COPY ACTIVE DATA` command. If you use the `ACTIVEDATAPOOLS` parameter, use the `COPY ACTIVE DATA` command to ensure that the active-data pools contain all active data of the primary storage pool.

SHRED

Specifies whether data is physically overwritten when it is deleted. This parameter is optional. You can specify an integer 0 - 10.

If you specify a value of zero, the server deletes the data from the database. However, the storage that is used to contain the data is not overwritten, and the data exists in storage until that storage is reused for other data. It might be possible to discover and reconstruct the data after it is deleted. Changing the value (for example, resetting it to 0) does not affect data that was deleted and is waiting to be overwritten.

If you specify a value greater than 0, the server deletes the data both logically and physically. The server overwrites the storage that is used to contain the data the specified number of times. This overwriting increases the difficulty of discovering and reconstructing the data after it is deleted.

To ensure that all copies of the data are shredded, specify a `SHRED` value greater than zero for the storage pool that is specified in the `NEXTSTGPOOL` parameter. Do not specify either the `COPYSTGPOOLS` or `ACTIVEDATAPOOLS`. Specifying relatively high values for the overwrite count generally improves the level of security, but might affect performance adversely.

Overwriting of deleted data is done asynchronously after the delete operation is complete. Therefore, the space that is occupied by the deleted data remains occupied for some time. The space is not available as free space for new data.

A `SHRED` value greater than zero cannot be used if the value of the `CACHE` parameter is `YES`. If you want to enable shredding for an existing storage pool for which caching is already enabled, you must change the value of the `CACHE` parameter to `NO`. Existing cached files remain in storage so that subsequent retrieval requests can be satisfied quickly. If space is needed to store new data, the existing cached files are erased so that the space they occupied can be used for the new data. The existing cached files are not shredded when they are erased.

Important: After an export operation finishes and identifies files for export, any change to the storage pool `SHRED` value is ignored. An export operation that is suspended retains the original `SHRED` value throughout the operation. You might want to consider canceling your export operation if changes to the storage pool `SHRED` value jeopardize the operation. You can reissue the export command after any needed cleanup.

Example: Update a random access storage pool to allow caching

Update the random access storage pool that is named `BACKUPPOOL` to allow caching when the server migrates files to the next storage pool.

```
update stgpool backuppool cache=yes
```

UPDATE STGPOOL (Update a primary sequential access pool)

Use this command to update a primary sequential access storage pool.

Restrictions:

1. You cannot use this command to change the data format for the storage pool.
2. If the value for `DATAFORMAT` is `NETAPPDUMP`, `CELERRADUMP`, or `NDMPDUMP`, you can modify only the following attributes:
 - o `DESCRIPTION`
 - o `ACCESS`
 - o `COLLOCATE`
 - o `MAXSCRATCH`
 - o `REUSEDELAY`

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax

```

+-MIGRation-+
'-All-----'

>----->
|          .-,'-----'. |
|          v          (1) (2) | |
|'-COPYSTGpools-----copypoolname-----+-'

>----->
|          (1) (2) |
|'-COPYContinue-----+Yes-+-----'
|          '-No--'

>----->
|          .-,'-----'. |
|          v          | |
|'-ACTIVEDATApools-----active-data_pool_name-+-'

>----->
|'-DEDUPlicate-----+No-----+-'
|          |          (3) |
|          '-Yes-----'

>-----><
|          (4) |
|'-IDENTIFYPRocess-----number-----'

```

Notes:

1. This parameter is not available for storage pools that use the data formats NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
2. This parameter is not available for CENTERA storage pools.
3. This parameter is valid only for storage pools that are defined with a FILE-type device class.
4. This parameter is only available if the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be updated.

DEScRiption

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. You can specify the following values:

READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool was specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. You can specify the following values:

NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer from 1 to 999999 terabytes, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

Scale factor	Meaning
K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

The client estimates the size of files that are sent to the server. The client estimate is used rather than the actual amount of data that is sent to the server. Client options, such as deduplication, compression, and encryption, can cause the actual amount of data that is sent to the server to be larger or smaller than the size estimate. For example, the compression of a file might be smaller in size than the estimate, thus sending less data than the estimate. Furthermore, a binary file might be larger in size after the compression processing, thus sending more data than the estimate.

When the physical size of the storage pool exceeds the MAXSIZE parameter, the following table shows where files are typically stored.

Table 1. The location of a file according to the file size and the pool that is specified

File size	Pool specified	Result
Exceeds the maximum size	No pool is specified as the next storage pool in the hierarchy	The server does not store the file
	A pool is specified as the next storage pool in the hierarchy	The server stores the file in the next storage pool that can accept the file size

Tip: If you also specify the NEXTstgpool parameter, define one storage pool in your hierarchy to have no limit on the maximum file size by specifying the MAXSize=NOLimit parameter. When you have at least one pool with no size limit, you ensure that no matter what its size, the server can store the file.

For multiple files that are sent in a single transaction, the server considers the size of the transaction to be the file size. If the total size of all files in the transaction is larger than the maximum size limit, the server does not store the files in the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage

pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

NEXTstgpool

Specifies a primary storage pool to which files are migrated. You cannot migrate data from a sequential access storage pool to a random access storage pool. This parameter is optional. The next storage pool must be a primary storage pool.

To remove an existing value, specify a null string ("").

If this storage pool does not have a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

When there is insufficient space available in the current storage pool, the NEXTSTGPOOL parameter for sequential access storage pools does not allow data to be stored into the next pool. In this case, the server issues a message and the transaction fails.

For next storage pools with a device type of FILE, the server completes a preliminary check to determine whether sufficient space is available. If space is not available, the server skips to the next storage pool in the hierarchy. If space is available, the server attempts to store data in that pool. However, it is possible that the storage operation might fail because, at the time the actual storage operation is attempted, the space is no longer available.

Restrictions:

- To ensure that you do not create a chain of storage pools that leads to an endless loop, specify at least one storage pool in the hierarchy with no value.
- If you specify a sequential-access pool as the next storage pool, the pool must be in either NATIVE or NONBLOCK data format.
- Do not specify a directory-container or cloud-container storage pool.
- Do not use this parameter to specify a storage pool for data migration.
- This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP

HIghmig

Specifies that the server starts migration when storage pool utilization reaches this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer 0 - 100.

When the storage pool exceeds the high migration threshold, the server can start migration of files by volume to the next storage pool defined for the storage pool. You can set the high migration threshold to 100 to prevent migration for the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

LOWmig

Specifies that the server stops migration when storage pool utilization is at or below this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer 0 - 99.

When the storage pool reaches the low migration threshold, the server does not start migration of files from another volume. You can set the low migration threshold to 0 to allow migration to empty the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

REClaim

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the IBM Spectrum Protect™ database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

AIX | **Windows** For storage pools that use a WORM device class, you can lower the value from the default of 100. Lowering the value allows the server to consolidate data onto fewer volumes when needed. Volumes that are emptied by reclamation can be checked out of the library, freeing slots for new volumes. Because the volumes are write-once, the volumes cannot be reused.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIMPRocess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999. You can specify one or more reclamation processes for each primary sequential-access storage pool.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Assuming that the RECLAIMSTGPOOL parameter is not specified or that the reclaim storage pool has the same device class as the storage pool that is being reclaimed, each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the two storage pools must have a mount limit of at least 16.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIMSTGpool

Specifies another primary storage pool as a target for reclaimed data from this storage pool. This parameter is optional. When the server reclaims volumes for the storage pool, unexpired data is moved from the volumes that are being reclaimed to the storage pool named with this parameter.

To remove an existing value, specify a null string ("").

A reclaim storage pool is most useful for a storage pool that has only one drive in its library. When you specify this parameter, the server moves all data from reclaimed volumes to the reclaim storage pool regardless of the number of drives in the library.

To move data from the reclaim storage pool back to the original storage pool, use the storage pool hierarchy. Specify the original storage pool as the next storage pool for the reclaim storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

COLlocate

Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required. Collocation can also impact the number of processes that are migrating disks to sequential pool.

You can specify one of the following options:

No

Specifies that collocation is disabled. During migration from disk, processes are created at a file space level.

GRoup

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.
- During migration from disk, the server creates migration processes at the collocation group level for grouped nodes, and at the node level for ungrouped nodes.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces that are named A, B, C, D, and E. File spaces A and B belong to a file space collocation group but C, D, and E do not. File spaces A and B are collocated by file space collocation group, while C, D, and E are collocated by node.

- During migration from disk, the server creates migration processes at the collocation group level for grouped file spaces.

Data is collocated on the least number of sequential access volumes.

NODe

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

For COLLOCATE=NODE, the server creates processes at the node level when you migrate data from disk.

FILESpace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

For COLLOCATE=FILESPACE, the server creates processes at the file space level when you migrate data from disk.

MAXSCRatch

Specifies the maximum number of scratch volumes that the server can request. This parameter is optional. You can specify an integer 0 - 100000000. By allowing the server to request scratch volumes, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the storage pool and the corresponding estimated capacity for the storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. When scratch volumes with the device type of FILE are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. The value 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

By specifying this parameter, you can ensure that the database can be restored to an earlier level and database references to files in the storage pool would still be valid.

OVFLOcation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the MOVE MEDIA command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. All files on a volume must be eligible for migration before the server selects the volume for migration. To calculate a value to compare to the specified MIGDELAY, the server counts the number of days that the file has been in the storage pool.

This parameter is optional. You can specify an integer 0 - 9999.

If you want the server to count the number of days that are based only on when a file was stored and not when it was retrieved, use the NORETRIEVEDATE server option.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue migration by migrating files that do not satisfy the migration delay time.

You can specify one of the following values:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that have not been stored in the storage pool for the number of days specified by the migration delay period.

No

Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files have been stored in the storage pool for the number of days specified by the migration delay period.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGProcess

Specifies the number of parallel processes to use for migrating the files from the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the migration, and the number of logical and physical drives that can be dedicated to the operation. To access a sequential-access volume, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other IBM Spectrum Protect and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the migration.

For example, suppose you want to simultaneously migrate the files from volumes in two primary sequential storage pools and that you want to specify three processes for each of the storage pools. The storage pools have the same device class. Assuming that the storage pool to which files are being migrated has the same device class as the storage pool from which files are being migrated, each process requires two mount points and, if the device type is not FILE, two drives. (One drive is for the input volume, and the other drive is for the output volume.) To run six migration processes simultaneously, you need a total of at least 12 mount points and 12 drives. The device class for the storage pools must have a mount limit of at least 12.

If the number of migration processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the migration processes will end. For information about specifying the MOUNTWAIT time, see DEFINE DEVCLASS (Define a device class).

The IBM Spectrum Protect server will start the specified number of migration processes regardless of the number of volumes that are eligible for migration. For example, if you specify ten migration processes and only six volumes are eligible for migration, the server will start ten processes and four of them will complete without processing a volume.

Note: When you specify this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool that is defined to the target storage pool.

AUTOCopy

Specifies when IBM Spectrum Protect completes simultaneous-write operations. This parameter affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If the AUTOCOPY option is set to `ALL` or `CLIENT`, and there is at least one storage pool that is listed in the `COPYSTGPOLS` or `ACTIVEDATAPOOLS` options, any client-side deduplication is disabled.

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the `COPYSTGPOLS` parameter. Active-data pools are specified using the `ACTIVEDATAPOOLS` parameter.

You can specify one of the following values:

`None`

Specifies that the simultaneous-write function is disabled.

`CLient`

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

`MIGRation`

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

`All`

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

`COPYSTGpools`

Specifies the names of copy storage pools where the server simultaneously writes data. You can specify a maximum of three copy pool names that are separated by commas. Spaces between the names of the copy pools are not allowed. To add or remove one or more copy storage pools, specify the pool name or names that you want to include in the updated list. For example, if the existing copy pool list includes `COPY1` and `COPY2` and you want to add `COPY3`, specify `COPYSTGPOLS=COPY1,COPY2,COPY3`. To remove all existing copy storage pools that are associated with the primary storage pool, specify a null string (""), for example, `COPYSTGPOLS=""`.

When you specify a value for the `COPYSTGPOLS` parameter, you can also specify a value for the `COPYCONTINUE` parameter. For more information, see the `COPYCONTINUE` parameter.

The combined total number of storage pools that are specified in the `COPYSTGPOLS` and `ACTIVEDATAPOOLS` parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the `COPYCONTINUE` value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to copy storage pools during the following operations:

- Back up and archive operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API
- Migration operations by IBM Spectrum Protect for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

Restrictions:

1. This parameter is available only to primary storage pools that use `NATIVE` or `NONBLOCK` data format. This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Simultaneous-write operations takes precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is accepted.
 3. The simultaneous-write function is not supported for NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.
 4. You cannot use the simultaneous-write function with CENTERA storage devices.

Attention: The function that is provided by the COPYSTGPOOLS parameter is not intended to replace the BACKUP STGPOOL command. If you use the COPYSTGPOOLS parameter, continue to use the BACKUP STGPOOL command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

COPYContinue

Specifies how the server reacts to a copy storage pool write failure for any of the copy storage pools that are listed in the COPYSTGPOOLS parameter. This parameter is optional. The default is YES. When you specify the COPYCONTINUE parameter, either a COPYSTGPOOLS list must exist or the COPYSTGPOOLS parameter must also be specified.

The COPYCONTINUE parameter has no effect on the simultaneous-write function during migration.

You can specify the following values:

Yes

If the COPYCONTINUE parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

No

If the COPYCONTINUE parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the COPYCONTINUE parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the COPYCONTINUE parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The ACTIVEDATAPOOLS parameter is optional. Spaces between the names of the active-data pools are not allowed.

The combined total number of storage pools that are specified in the COPYSSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data.

The server can write data simultaneously to active-data pools only during backup operations by IBM Spectrum Protect backup-archive clients or application clients that use the IBM Spectrum Protect API.

Restrictions:

1. This parameter is available only to primary storage pools that use NATIVE or NONBLOCK data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP

- o CELERRADUMP
 - o NDMPDUMP
2. Writing data simultaneously to active-data pools is not supported when the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is accepted.
 3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the TOCDESTINATION in the copy group of the management class has active-data pools defined, the active-data pools are ignored and the data is stored into the primary storage pool only.
 4. You cannot use the simultaneous-write function with CENTERA storage devices.
 5. Data being imported cannot be stored in active-data pools. After an import operation, use the COPY ACTIVE DATA command to store the imported data in an active-data pool.

Attention: The function that is provided by the ACTIVE DATA POOLS parameter is not intended to replace the COPY ACTIVE DATA command. If you use the ACTIVE DATA POOLS parameter, use the COPY ACTIVE DATA command to ensure that the active-data pools contain all active data of the primary storage pool.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE device class.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a device class associated with the FILE device type. Enter a value 1 - 50. Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Update the primary sequential storage pool's mountable scratch volumes

Update the primary sequential storage pool that is named TAPEPOOL1 to allow as many as 10 scratch volumes to be mounted.

```
update stgpool tapepool1 maxscratch=10
```

UPDATE STGPOOL (Update a copy sequential access storage pool)

Use this command to update a copy sequential access storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax

```
>>-UPDate STGpool--pool_name--+-+-----+-----+-----+-----+----->
                                     '-DEScRiption----description-'
>--+-----+-----+-----+-----+----->
  '-ACCess----+-READWrite----+'
      +-READOnly-----+
      '-UNAVailable-'
>--+-----+-----+-----+-----+----->
  '-COLlocate----+-No-----+'   '-REClaim----percent-'
      +-GRoup-----+
      +-NODE-----+
      '-Filespace-'
>--+-----+-----+-----+-----+----->
  '-RECLAIMProcess----number-'
>--+-----+-----+-----+-----+----->
```

```

'-OFFSITERECLAIMLimit-----+NOLimit+-'
      '-number--'

>-----+-----+-----+----->
'-MAXSCRatch-----number-' '-REUsedelay-----days-'

>-----+-----+-----+----->
'-OVFLocation-----location-' '-CRCData-----+Yes+-'
      '-No--'

>-----+-----+-----+----->
'-DEDuplicate-----+No-----+-'
      | (1) |
      '-Yes-----'

>-----+-----+-----+-----><
      | (2) |
      '-IDENTIFYPRocess-----number-----'

```

Notes:

1. This parameter is valid only for storage pools that are defined with a FILE-type device class.
2. This parameter is only available if the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the copy storage pool to be updated.

DEScRiption

Specifies a description of the copy storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how client nodes and server processes (such as reclamation) can access files in the copy storage pool. This parameter is optional. You can specify the following values:

READWrite

Specifies that files can be read from and written to the volumes in the copy storage pool.

READOnly

Specifies that client nodes can read only files that are stored on the volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are allowed to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

UNAVailable

Specifies that client nodes cannot access files that are stored on volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are allowed to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

COLlocate

Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required.

You can specify one of the following options:

No

Specifies that collocation is disabled.

GROUp

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces named A, B, C, D, and E. File spaces A and B belong to a file space collocation group but C, D, and E do not. File spaces A and B are collocated by file space collocation group, while C, D, and E are collocated by node.

Data is collocated on the least amount of sequential access volumes.

NODE

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

FILEspace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

REClaim

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume.

Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the IBM Spectrum Protect™ database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100. The value 100 means that reclamation is not completed.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

If you change the value from the default of 100, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the active files on the reclaimable volume from a primary or copy storage pool that is onsite. The process then writes these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location. However, the files can be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

RECLAIMProcess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for each storage pool must have a mount limit of at least eight.

You can specify one or more reclamation processes for each copy storage pool. You can specify multiple concurrent reclamation processes for a single copy storage pool, which makes better use of your available tape drives or FILE volumes. If multiple concurrent processing is not necessary, specify a value of 1 for the RECLAIMPROCESS parameter.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes that space is reclaimed from during reclamation for this storage pool. This parameter is optional. You can specify the following values:

NOLimit

Specifies that you want to reclaim the space in all of your offsite volumes.

number

Specifies the number of offsite volumes to reclaim space from. You can specify an integer 0 - 99999. A value of zero means that none of the offsite volumes are reclaimed.

Tip:

To determine the value for the OFFSITERECLAIMLIMIT, use the statistical information in the message that is issued at the end of the offsite volume reclamation operation. The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose a copy storage pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 will be reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 will be reclaimed.

MAXSCRatch

Specifies the maximum number of scratch volumes that the server can request for this storage pool. This parameter is optional. You can specify an integer 0 - 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the copy storage pool and the corresponding estimated capacity for the copy storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the copy storage pool until the access mode is changed. An administrator can query the server for empty, offsite scratch volumes and return them to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The IBM Spectrum Protect server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. A value of 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to files in the copy storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the MOVE MEDIA command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value 1 - 50.

Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active.

Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes

and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Update a copy storage pool to a 30-day volume reuse and to collocate files by client node

Update the copy storage pool that is named TAPEPOOL2 to change the delay for volume reuse to 30 days and to collocate files by client node.

```
update stgpool tapepool2 reusedelay=30 collocate=node
```

Related reference:

SET DRMDBBACKUPEXPIREDDAYS (Specify DB backup series expiration)

UPDATE STGPOOL (Update an active-data sequential access)

Use this command to update an active-data pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax

```
>>-UPDate STGpool--pool_name--+-+-----+-----+-----+-----+----->
                                     '-DEScRiption-----description-'
>--+-----+-----+-----+-----+----->
  '-ACCess-----+READWrite---+'
      +-READOnly-----+
      '-UNAVailable-'
>--+-----+-----+-----+-----+----->
  '-COLlocate-----+No-----+'   '-RECLaim-----percent-'
      +-GRoup-----+
      +-NODE-----+
      '-Filespace-'
>--+-----+-----+-----+-----+----->
  '-RECLAIMPRocess-----number-'
>--+-----+-----+-----+-----+----->
  '-OFFSITERECLAIMLimit-----+NOLimit+-+'
                                     '-number--'
>--+-----+-----+-----+-----+----->
  '-MAXSCRatch-----number-'   '-REUsedelay-----days-'
>--+-----+-----+-----+-----+----->
  '-OVFLocation-----location-'   '-CRCDATA-----+Yes+-+'
                                     '-No--'
>--+-----+-----+-----+-----+----->
  '-DEDUPlicate-----+No-----+'
      |           (1) |
      '-Yes-----'
>--+-----+-----+-----+-----+-----><
  |           (2) |
  '-IDENTIFYPRocess-----number-----'
```

Notes:

1. This parameter is valid only for storage pools that are defined with a FILE-type device class.
2. This parameter is only available if the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the active-data pool to be updated.

DEscription

Specifies a description of the active-data pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

ACCess

Specifies how client nodes and server processes (such as reclamation) can access files in the active-data pool. This parameter is optional. You can specify the following values:

READWrite

Specifies that files can be read from and written to the volumes in the active-data pool.

READOnly

Specifies that client nodes can read only files that are stored on the volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore active versions of backup files to primary storage pools. However, no new writes are allowed to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

UNAVailable

Specifies that client nodes cannot access files that are stored on volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore active versions of backup files to primary storage pools. However, no new writes are allowed to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

COLlocate

Specifies whether the server attempts to keep data, which is stored on as few volumes as possible, that belong to one of the following candidates:

- A single client node
- A group of file spaces
- A group of client nodes
- A client file space

This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time that is needed to collocate files for storing and the number of volumes required.

You can specify one of the following options:

No

Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes or file spaces. For collocation groups, the server attempts to put data for nodes or file spaces that belong to the same collocation group on as few volumes as possible.

If you specify COLLOCATE=GROUP but do not define any collocation groups, or if you do not add nodes or file spaces to a collocation group, data is collocated by node. Consider tape usage when you organize client nodes or file spaces into collocation groups.

For example, if a tape-based storage pool consists of data from nodes and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates the data by group for grouped nodes. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes that are associated with a group.
- Collocates the data by node for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

If a tape-based storage pool consists of data from grouped file spaces and you specify COLLOCATE=GROUP, the server completes the following actions:

- Collocates by group, the data for grouped file spaces only. Whenever possible, the server collocates data that belongs to a group of file spaces on a single tape or on as few tapes as possible. Data for a single file space can also be spread across several tapes that are associated with a group.
- Collocates the data by node (for file spaces that are not explicitly defined to a file space collocation group). For example, node1 has file spaces named A, B, C, D, and E. File spaces A and B belong to a filespace collocation group but C, D, and E do not. File spaces A and B are collocated by filespace collocation group, while C, D, and E are collocated by node.

Data is collocated on the least amount of sequential access volumes.

NODE

Specifies that collocation is enabled at the client node level. For collocation groups, the server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not try to collocate those file spaces. For compatibility with an earlier version, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data is collocated by node.

FILESPACE

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to place data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to place data for different file spaces on different volumes.

RECLAIM

Specifies when the server reclaims a volume, which is based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space that is occupied by files that are expired or deleted from the IBM Spectrum Protect™ database.

Reclamation makes the fragmented space and space occupied by inactive backup files on volumes usable again by moving any remaining unexpired files and active backup files from one volume to another volume. This action makes the original volume available for reuse. This parameter is optional. You can specify an integer 1 - 100. The value 100 means that reclamation is not completed.

The server determines that the volume is a candidate for reclamation if the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool.

If you change the value from the default of 60, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When an active-data pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the active files on the reclaimable volume from a primary or active-data pool that is onsite. The process then writes these files to an available volume in the original active-data pool. Effectively, these files are moved back to the onsite location. However, the files can be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with active-data pools.

RECLAIMPROCESS

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value 1 - 999.

When you calculate the value for this parameter, consider the following resources that are required for reclamation processing:

- The number of sequential storage pools.
- The number of logical and physical drives that can be dedicated to the operation.

To access sequential volumes, IBM Spectrum Protect uses a mount point and, if the device type is not FILE, a physical drive.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for each storage pool must have a mount limit of at least eight.

You can specify one or more reclamation processes for each active-data pool. You can specify multiple concurrent reclamation processes for a single active-data pool, which makes better use of your available tape drives or FILE volumes. If multiple concurrent processing is not necessary, specify a value of 1 for the RECLAIMPROCESS parameter.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes that space is reclaimed from during reclamation for this storage pool. This parameter is optional. You can specify the following values:

NOLimit

Specifies that you want to reclaim the space in all of your offsite volumes.

number

Specifies the number of offsite volumes to reclaim space from. You can specify an integer 0 - 99999. A value of zero means that none of the offsite volumes are reclaimed.

Tip:

To determine the value for the OFFSITERECLAIMLIMIT, use the statistical information in the message that is issued at the end of the offsite volume reclamation operation. The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose an active-data pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes are reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 are reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 is reclaimed.

MAXSCRatch

Specifies the maximum number of scratch volumes that the server can request for this storage pool. This parameter is optional. You can specify an integer 0 - 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value that is specified for this parameter is used to estimate the total number of volumes available in the active-data pool and the corresponding estimated capacity for the active-data pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the active-data pool until the access mode is changed. An administrator can query the server for empty, offsite scratch volumes and return them to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, IBM Spectrum Protect marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition is not reached. The IBM Spectrum Protect server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer 0 - 9999. A value of 0 means that a

volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

Tip: Use this parameter to ensure that when you restore the database to an earlier level, database references to files in the active-data pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days that are specified for this parameter must be the same as the number specified for the SET DRMDBBACKUPEXPIREDAYS command.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the MOVE MEDIA command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command, you can continually ensure the integrity of data that is stored in your storage hierarchy. You can specify the following values:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because more processing is required to calculate and compare CRC values between the storage pool and the server.

No

Specifies that data is stored without CRC information.

Tip:

For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the LBPROTECT parameter on the DEFINE DEVCLASS and UPDATE DEVCLASS commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM® LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C and T10000D drives.

DEDuplicate

Specifies whether the data that is stored in this storage pool is deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class.

IDENTIFYProcess

Specifies the number of parallel processes to use for server-side data deduplication. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value 1 - 50.

Remember: Data deduplication processes can be either active or idle. Processes that are working on files are active.

Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the QUERY PROCESS command for data deduplication includes the total number of bytes and files that have been processed since the process first started. For example, if a data deduplication process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of data-deduplication processes for the storage pool is changed to a value less than the number currently specified.

Example: Update an active data pool

Update the active-data pool that is named TAPEPOOL2 to change the delay for volume reuse to 30 days and to collocate files by client node.

```
update stgpool tapepool3 reusedelay=30 collocate=node
```

Related reference:

SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)

AIX

Linux

Windows

UPDATE STGPOOLDIRECTORY (Update a storage pool directory)

Use this command to update a storage pool directory.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```
>>-UPDate STGPOOLDIRectory--pool_name--directory----->
                                     .-MAXPRocess----4-----.
>----ACCess---+--READWrite---+-----+-----+----->
          +--READOnly----+   '-MAXProcess----number-'
          +--DEStroyed---+
          '-UNAVailable-'

      .-Wait-----No-----.
>--+-----+-----><
      '-Wait-----+--No---+'
          '-Yes-'
```

Parameters

pool_name (Required)

Specifies the storage pool that contains the directory to update. This parameter is required.

directory (Required)

Specifies a file system directory of the storage pool. This parameter is required.

ACCess (Required)

Specifies how client nodes and server processes can access files in the storage pool directory. This parameter is required. The following values are possible:

READWrite

Specifies that files can be read from and written to the storage pool directory.

READOnly

Specifies that files can be read from the storage pool directory.

DEStroyed

Specifies that files are permanently damaged and must be destroyed in the storage pool directory. Use this access mode to indicate that an entire storage pool directory must be recovered.

Tips:

- Mark storage pool directories as `DESTROYED` before you complete data recovery. When the storage pool directory is marked as destroyed, you can recover data extents on the target replication server.
- Use the `MAXPROCESS` parameter to specify the number of parallel processes that you can use to update a storage pool directory.

UNAVailable

Specifies that files cannot be accessed on the storage pool directory in the storage pool.

MAXPRocess

Specifies the maximum number of parallel processes to use for updating a storage pool directory. This parameter is optional. You can enter a value in the range 1 - 99. The default value is 4.

Restriction: You can use this parameter only when you specify the `ACCESS=DESTROYED` parameter.

When you specify the `ACCESS=DESTROYED` parameter, each container in the storage pool directory is updated by one process. If the maximum number of parallel processes is larger than or equal to the number of containers that must be updated, only one process is created for each container. If the number of containers exceeds the value of the `MAXPROCESS` parameter, the command waits for the child processes to finish before any new processes can begin.

Wait

This optional parameter specifies whether to wait for the IBM Spectrum Protect™ server to complete processing this command in the foreground. The default is NO. You can specify the following values:

No

The server processes this command in the background and you can continue with other tasks while the command is processing. Messages that are related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must complete processing before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Update a storage pool directory to destroy it

Update a storage pool directory that is named DIR1 in storage pool POOL1 to mark it as destroyed.

```
update stgpooldirectory pool1 dir1 access=destroyed
```

Example: Update a storage pool directory to destroy it in a cloud-container storage pool

Update a storage pool directory that is named DIR3 in cloud-container storage pool CLOUDLOCALDISK1 to mark it as destroyed.

```
update stgpooldirectory cloudlocaldisk1 dir3 access=destroyed
```

Example: Update a storage pool directory to make it unavailable

When the storage pool directory is unavailable, the server does not read or write data to the directory. To update the access mode to unavailable for a storage pool directory, `dir1`, in a storage pool that is named `pool1`, issue the following command:

```
update stgpooldirectory pool1 dir1 access=unavailable
```

Table 1. Commands related to UPDATE STGPOOLDIRECTORY

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DEFINE STGPOOLDIRECTORY	Defines a storage pool directory to a directory-container or cloud-container storage pool.
DELETE STGPOOLDIRECTORY	Deletes a storage pool directory from a directory-container or cloud-container storage pool.
QUERY STGPOOLDIRECTORY	Displays information about storage pool directories.

UPDATE STGRULE (Update a storage rule)

Use this command to update a storage rule.

The UPDATE STGRULE command takes several forms. The syntax and parameters for each form are defined separately.

Table 1. Commands related to UPDATE STGRULE

Command	Description
DEFINE STGRULE (auditing)	Defines a storage rule for auditing storage pools.
DEFINE STGRULE (data deduplication statistics)	Defines a storage rule for generating data deduplication statistics.
DEFINE STGRULE (reclaiming)	Defines a storage rule for reclaiming cloud-container storage pools.
DEFINE STGRULE (tiering)	Defines a storage rule for tiering.
DELETE STGRULE	Deletes storage rules.

Command	Description
QUERY STGRULE	Displays storage rule information.

- UPDATE STGRULE (Update a rule for auditing a storage pool)
Use this command to update a rule that schedules audit operations for a storage pool.
- UPDATE STGRULE (Update a storage rule for generating data deduplication statistics)
Use this command to update a storage rule for generating data deduplication statistics.
- UPDATE STGRULE (Update a storage rule for reclaiming cloud containers)
Use this command to update a storage rule for reclaiming space in cloud-container storage pools.
- UPDATE STGRULE (Update a storage rule for tiering)
Use this command to update a storage rule for one or more storage pools. The storage rule schedules tiering between container storage pools. You can update one or more storage rules for a container storage pool.

UPDATE STGRULE (Update a rule for auditing a storage pool)

Use this command to update a rule that schedules audit operations for a storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```

                .-DELAY-----7-----.
>>-Update STGRULE--rule_name--+-----+----->
                '-DELAY-----delay-'

    .-AUDITType-----Extent-.    .-AUDITLevel-----5-----.
>>-+-----+-----+-----+----->
                '-AUDITLevel-----+1--+-'
                                '-5-'

    .-STARTTime-----current_time-.    .-ACTIVE-----Yes-----.
>>-+-----+-----+-----+----->
    '-STARTTime-----time-----'    '-ACTIVE-----+No--+-'
                                        '-Yes-'

>>-+-----+-----+-----+----->>
    '-DESCription-----description-'

```

Parameters

rule_name (Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

DELAY

Specifies the interval, in days, between audit operations. This parameter is optional. The default value is 7 days. You can specify an integer in the range 1 - 9999.

AUDITType

Specifies the audit type. This parameter is optional. You can specify the following value:

Extent

Specifies that only extents are audited. This is the default value.

Restriction: In IBM Spectrum Protect™ Version 8.1.5, you can use the audit storage rule only to audit extents. Objects are not audited.

AUDITLevel

Specifies the level of the audit. This parameter is optional. The following values are possible:

1

Specifies a minimal audit operation of the extents in the storage pool.

5

Specifies a full audit operation of the extents in the storage pool. This is the default value.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. The default is the current time. This parameter is optional.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether storage rule processing occurs. This parameter is optional. The following values are possible:

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time. This is the default value.

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

DEscription

Specifies a description of the storage rule. This parameter is optional. The maximum length of the description is 255 characters. If the description includes spaces, enclose the description in quotation marks.

Update a rule for an extent-level audit operation

Update a storage rule, AUDITACCOUNTING, to schedule a full, extent-level audit of data starting at 3 AM. The audit operation takes place every 14 days:

```
update stgrule auditaccounting delay=14 auditlevel=5 starttime=03:00:00
```

Related commands

Table 1. Commands related to UPDATE STGRULE

Command	Description
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.
UPDATE STGRULE (auditing)	Updates a storage rule for auditing storage pools.

UPDATE STGRULE (Update a storage rule for generating data deduplication statistics)

Use this command to update a storage rule for generating data deduplication statistics.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```
>>-UPDate STGRULE--rule_name--+-----+----->
                                     '-DELAY---delay-'
>--+-----+-----+-----+----->
   '-MAXPRocess---number-'   '-STARTTime---time-'
>--+-----+-----+-----+----->
```

```

'-ACTIVE-----+No--+-'
                '-Yes-'

>-----+----->
|               .-,------. |
|               v               |
|'-NODEList-----+node_name-----+-'
|                   '-node_group_name-'

>-----+----->
'-NAMEType-----+SERVER--+-'
                +-UNICODE-+
                '-FSID-----'

>-----+----->
|               .-,------. |
|               v               |
|'-FSLIST-----+file_space_name--+-'
|                   +-,-----+
|                   '-fsid-----'

>-----+----->
'-CODEType-----+UNICODE-----+'
                +-NONUNICODE-+
                '-BOTH-----'

>-----+----->>
'-DESCRIPTION-----description-'

```

Parameters

rule_name (Required)

Specifies the name of the storage rule. The name must be unique, and the maximum length is 30 characters.

DELAY

Specifies the number of days to wait before the statistics are generated. You can specify an integer in the range 0 - 9999.

MAXProcess

Specifies the maximum number of parallel processes to collect statistics for each storage pool that is specified. This parameter is optional. You can enter a value in the range 1 - 99. For example, if you have 4 storage pools and you specify a value of 8, 32 processes are started.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether storage rule processing occurs. This parameter is optional. The following values are possible:

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

NODEList

Specifies the name of the client node or defined group of client nodes for which data deduplication statistics are collected. You can also specify a combination of client node names and client-node group names. To specify multiple client node names or client-node group names, separate the names with commas with no intervening spaces. You can use wildcard

characters with client node names but not with client-node group names. The specified value can have a maximum of 1024 characters. If you enter an asterisk (*), information is shown for all client nodes. This parameter is optional.

NAMEType

Specifies how you want the server to interpret the file space names that you enter. Use this parameter when IBM Spectrum Protect™ clients have file spaces that are in Unicode format, and are on Windows, NetWare, or Macintosh OS X operating systems. This parameter is optional.

This parameter is required if you specify a node name and a file space name or FSID.

Restriction: When you specify this parameter, the file space name cannot contain an asterisk.

You can specify one of the following values:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space name that is entered from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server's code page.

Tip: Conversion can fail if the string includes characters that are not available in the server code page, or if the server cannot access system conversion routines.

FSID

The server interprets the file space names as their FSIDs.

FSLIST

Specifies the names of one or more file spaces for which data deduplication statistics are collected. This parameter is optional. You can use wildcard characters to specify this name. The specified value can have a maximum of 1024 characters. You can specify one of the following values:

*

Specify an asterisk (*) to show information for all file spaces or IDs.

filespace_name

Specifies the name of the file space. You can specify more than one file space by separating the names with commas and no intervening spaces.

FSID

Specifies the name of a file space identifier. This parameter is valid for clients with file spaces that are in Unicode format. Specify more than one file space by separating the names with commas and no intervening spaces.

For clients with file spaces that are in Unicode format, you can enter either a file space name or an FSID. If you enter a file space name, the server might have to convert the file space name that you enter. For example, the server might have to convert the name that you enter from the server's code page to Unicode.

Restrictions: The following restrictions apply to file space names and FSIDs:

- You must specify a node name if you specify a file space name.
- Do not specify both file space names and FSIDs on the same command.

CODEType

Specifies what type of file spaces to include in the record. Use this parameter only when you enter an asterisk to display information about all file spaces. This parameter is optional. Specify one of the following values:

UNICODE

Include file spaces that are in Unicode format.

NONUNICODE

Include file spaces that are not in Unicode format.

BOTH

Include file spaces regardless of code page type.

DESCRIPTION

Specifies a description of the storage rule. This parameter is optional.

Update a rule to generate data deduplication statistics

Update a storage rule that is named MYSTAT1 to generate data deduplication statistics. Limit the scope to the node that is named NODE1:

```
update stgrule mystat1 nodelist=node1
```


Related commands

Table 1. Commands related to UPDATE STGRULE

Command	Description
DEFINE STGRULE (data deduplication statistics)	Defines a storage rule for generating data deduplication statistics.
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.

UPDATE STGRULE (Update a storage rule for reclaiming cloud containers)

Use this command to update a storage rule for reclaiming space in cloud-container storage pools.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Restriction: You can configure a cloud reclamation rule for a storage pool only on a Microsoft Azure cloud computing system or on a cloud computing system with the Simple Storage Service (S3) protocol.

Syntax

```
>>-Update STGRULE--rule_name--+----->
                               '-PCTUnused----percentage-'
>--+----->
  '-MAXProcess----number-'  '-DURation----minutes-'
>--+----->
  '-STARTTime----time-'  '-ACTIVE----+No--+-'
                               '-Yes-'
>--+----->>
  '-DESCRiption----description-'
```

Parameters

rule_name (Required)

Specifies the name of the storage rule.

PCTUnused

Specifies the percentage of the cloud container that is no longer in use. This parameter is optional. After unused space reaches the specified value, the cloud container is reclaimed. You can specify an integer in the range 50 - 99.

MAXProcess

Specifies the maximum number of parallel processes for each reclamation operation. This parameter is optional. You can specify an integer in the range 1 - 99.

DURation

Specifies the maximum number of minutes that the storage rule runs before it is automatically canceled. You can specify a number in the range 60 - 1440. If you do not specify a value, the duration is not updated. You can specify the NOLIMIT parameter to allow the rule to run to completion. This parameter is optional.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

You can specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00

Value	Description	Example
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether the storage rule processing occurs. This parameter is optional. The following values are possible:

No

Specifies that the storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the storage rule is active. The storage rule is processed at the scheduled time.

DEScRiption

Specifies a description of the storage rule. This parameter is optional.

Update a rule to reclaim cloud containers

Update a storage rule that is named RECLAIMRULE to reclaim cloud containers that no longer use 60 percent of their space. Specify a start time of 23:30:00:

```
update stgrule reclaimrule pctunused=60 starttime=23:30:00
```

Related commands

Table 1. Commands related to UPDATE STGRULE

Command	Description
DEFINE STGRULE (reclaiming)	Defines a storage rule for reclaiming cloud-container storage pools.
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.

UPDATE STGRULE (Update a storage rule for tiering)

Use this command to update a storage rule for one or more storage pools. The storage rule schedules tiering between container storage pools. You can update one or more storage rules for a container storage pool.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege.

Syntax

```
>>-UPDate STGRULE--rule_name----->
>--+-----+-----+-----+----->
| .,----- . | '-TIERDelay---delay-'
| v           | |
|'---SRCPools---source_pool+--'
>--+-----+-----+-----+----->
' -MAXProcess---number-' ' -DURation---+minutes+-'
                               '-NOLimit-'
>--+-----+-----+-----+----->
' -STARTTime---time-' ' -ACTIVE---+No+--'
                               '-Yes-'
>--+-----+-----+-----+-----><
' -DEScRiption---description-'
```

Parameters

rule_name(Required)

Specifies the name of the storage rule. The maximum length of the name is 30 characters.

SRCPools

Specifies the name of one or more directory-container storage pools from which objects are tiered to the target storage pool. To specify multiple storage pools, separate the names with commas with no intervening spaces.

TIERDelay

Specifies the number of days to wait before the storage rule tiers objects to the next storage pool. You can specify an integer in the range 0 - 9999. The parameter value applies to all files in the storage pool.

MAXProcess

Specifies the maximum number of parallel processes to complete the storage rule for each source storage pool that is specified. This parameter is optional. Enter a value in the range 1 - 99. For example, if you have 4 source storage pools and you specify the default value of 8 for this parameter, 32 processes are started.

DURation

Specifies the maximum number of minutes that the storage rule runs before it is automatically canceled. You can specify a number in the range 60 - 1440. If you specify a value of NOLimit, the storage rule runs until it is completed. This parameter is optional.

STARTTime

Specifies the time for the beginning of the window in which the storage rule is first processed. This parameter is optional. The storage rule runs daily within 5 minutes following the specified time.

Specify one of the following values:

Value	Description	Example
HH:MM:SS	A specific time.	23:30:08
NOW	The current time.	NOW
NOW+HH:MM or +HH:MM	The current time plus the specified number of hours and minutes.	NOW+02:00 or +02:00
NOW-HH:MM or -HH:MM	The current time minus the specified number of hours and minutes.	NOW-02:00 or -02:00

ACTIVE

Specifies whether the storage rule processing occurs. This parameter is optional. The following values are possible:

No

Specifies that the defined storage rule is inactive. The storage rule is not processed at the scheduled time.

Yes

Specifies that the defined storage rule is active. The storage rule is processed at the scheduled time.

DEScRiption

Specifies a description of the storage rule. This parameter is optional.

Update a storage rule

Update a storage rule that is named tieraction to move data from directory-container storage pools dirpool1 and dirpool2 to the cloud-container storage pool cloudpool1. Specify a start time of 23:30:08 hours and a maximum of 16 processes:

```
update stgrule tieraction srcpools=dirpool1,dirpool2
maxprocess=16 starttime=23:30:08
```

Related commands

Table 1. Commands related to UPDATE STGRULE

Command	Description
DEFINE STGRULE (tiering)	Defines a storage rule for tiering.
DELETE STGRULE	Deletes storage rules.
QUERY STGRULE	Displays storage rule information.

UPDATE VIRTUALFSMAPPING (Update a virtual file space mapping)

Use this command to update a virtual file space mapping definition.

Restriction: You cannot use the UPDATE VIRTUALFSMAPPING command to update a virtual file space mapping for an EMC Celerra or EMC VNX NAS device. You must use the DEFINE VIRTUALFSMAPPING command.

The NAS device needs an associated data mover definition because when the server updates a virtual file space mapping, the server contacts the NAS device to validate the virtual file system and file system name.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned

Syntax

```
>>-UPDate VIRTUALFSMapping--node_name--virtual_filespace_name--->
>--+-----+----->
  '-FILESystem-----new_file_system_name-'
>--+-----+-----><
  |                                     .-NAMEType-----SERVER-----|
  '-PATH-----new_path_name--+-----+-'
                                     '-NAMEType-----+SERVER-----+-'
                                     '-HEXadecimal-'
```

Parameters

node_name (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

virtual_filespace_name (Required)

Specifies the virtual file space mapping to update. You cannot use wildcard characters or specify a list of names.

FILESystem

Specifies the new name of the file system in which the path is located. The file system name must exist on the specified NAS node. The file system name cannot contain wildcard characters. The file system name should only be modified when the file system name is modified on the NAS device or, for example, the directory is moved to a different file system. This parameter is optional.

PATH

Specifies the new path from the root of the file system to the directory. The path can only reference a directory. This should only be modified when the path on the NAS device has changed; for example, the directory is moved to a different path. The maximum length of the path is 1024 characters. The path name is case sensitive. This parameter is optional.

NAMEType

Specifies how the server should interpret the path name specified. Specify this parameter only if you specify a path. This parameter is useful when a path contains characters that are not part of the code page on which the server is running. The default value is SERVER.

Possible values are:

SERVER

The code page in which the server is running is used to interpret the path.

HEXadecimal

The server interprets the path that you enter as the hexadecimal representation of the path. This option should be used when a path contains characters that cannot be entered. For example, this could occur if the NAS file system is set to a language different from the one in which the server is running.

Example: Modify the path of a virtual file space mapping

Update the virtual file space mapping named /mikeshomedir for the NAS node NAS1 by modifying the path.

```
update virtualfsmapping nas1 /mikeshomedir path=/new/home/mike
```

Related commands

Table 1. Commands related to UPDATE VIRTUALFSMAPPING

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
QUERY VIRTUALFSMAPPING	Query a virtual file space mapping.

UPDATE VOLHISTORY (Update sequential volume history information)

Use this command to update volume history information for a volume produced by a database backup or an export operation. This command does not apply to storage pool volumes.

Use the UPDATE BACKUPSET command to update specified backup set volume information in the volume history file. Do not use this UPDATE VOLHISTORY command to update backup set volume information in the volume history file.

Privilege class

You must have system privilege or unrestricted storage privilege to issue this command.

Syntax

```
>>-UPDate VOLHistory--volume_name----->
>--DEVclass---device_class_name--+-----+---->
                                '-LLocation-----location-'
>--+-----+-----><
  '-ORMState-----+Mountable-----+'
    +-NOTMountable-----+
    +-COUrier-----+
    +-VAult-----+
    '-COURIERRetrieve-'
```

Parameters

volume_name (Required)

Specifies the volume name. The volume must have been used for a database backup or an export operation.

DEVclass (Required)

Specifies the name of the device class for the volume.

LOcation

Specifies the volume location. This parameter is required if the ORMSTATE parameter is not specified. The maximum text length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

Tip: The UPDATE VOLHISTORY command supports updates to the location information and ORMSTATE for snapshot database backup volumes.

ORMState

Specifies a change to the state of a database backup volume. This parameter is required if the LOCATION parameter is not specified. This parameter is only supported for systems licensed with Disaster Recovery Manager. Possible states are:

- MOnutable
The volume contains valid data and is accessible for on-site processing.
- NOTMOnutable
The volume is on-site, contains valid data, and is not accessible for on-site processing.
- COUrier
The volume is being moved off-site.
- VAult
The volume is off-site, contains valid data, and is not accessible for on-site processing.
- COURIERRetrieve
The volume is being moved on-site.

Example: Update the location of a volume used for database backup

Update the location of a volume used for database backup, BACKUP1, to show that it has been moved to an off-site location.

```
update volhistory backup1 devclass=tapebkup
location="700 w. magee rd."
```

Related commands

Table 1. Commands related to UPDATE VOLHISTORY

Command	Description
BACKUP VOLHISTORY	Records volume history information in external files.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
MOVE DRMEDIA	Moves DRM media onsite and offsite.
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.

UPDATE VOLUME (Change a storage pool volume)

Use this command to change the access mode for one or more volumes in storage pools.

You can correct an error condition that is associated with a volume by updating the volume to an access mode of READWRITE. You can also use this command to change the location information for one or more volumes in sequential access storage pools.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```

(1)
>>-UPDate Volume-----volume_name----->
>--+-----+----->
  '-ACCess---+--READWrite-----+'
      +-READOnly-----+
      +-UNAVailable---+
      |           (2) |
      +-DESTROYed-----+
      |           (3) |
      '-OFFsite-----'
>--+-----+----->
  |           (4) |
  '-LOcation-----location-'

```

```

.-WHERESTGpool-----*-----
>-----+-----+----->
'-WHERESTGpool-----pool_name-'

.-WHEREDEVclass-----*-----
>-----+-----+----->
'-WHEREDEVclass-----device_class_name-'

>-----+-----+----->
|          .-,------. |
|          V          | |
'-WHEREACCess-----+READWrite---+--+'
|          +READOnly----+
|          +-UNAVailable--+
|          +-OFFsite-----+
|          '-DESTroyed---'

>-----+-----+----->
|          .-,------. |
|          V          | |
'-WHEREStatus-----+ONline---+--+'
|          +OFFline--+
|          +-EMPTy---+
|          +-PENding--+
|          +-FILLing--+
|          '-FULL----'

.-Preview-----No-----
>-----+-----+-----><
'-Preview-----+No---+'
|          '-Yes-'

```

Notes:

1. You must update at least one attribute (ACCESS or LOCATION).
2. This value is valid only for volumes in primary storage pools.
3. This value is valid only for volumes in copy, container-copy, and active-data storage pools.
4. This parameter is valid only for volumes in sequential access storage pools.

Parameters

volume_name (Required)

Specifies the storage pool volume to update. You can use wildcard characters to specify names.

ACCess

Specifies how client nodes and server processes (such as migration) can access files in the storage pool volume. This parameter is optional. Possible values are:

READWrite

Specifies that client nodes and server processes can read from and write to files stored on the volume.

If the volume that is being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

READOnly

Specifies that client nodes and server processes can only read files that are stored on the volume.

If the volume that is being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

UNAVailable

Specifies that neither client nodes nor server processes can access files that are stored on the volume.

Before making a random access volume unavailable, you must vary the volume offline. After you make a random access volume unavailable, you cannot vary the volume online.

If you make a sequential access volume unavailable, the server does not attempt to mount the volume.

If the volume that is being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

DESTROYED

Specifies that a primary storage pool volume has been permanently damaged. Neither client nodes nor server processes can access files that are stored on the volume. Use this access mode to indicate an entire volume that needs to be restored by using the RESTORE STGPOOL command. After all files on a destroyed volume are restored to other volumes, the server automatically deletes the destroyed volume from the database.

Only volumes in primary storage pools can be updated to DESTROYED.

Before you update a random access volume to DESTROYED access, you must vary the volume offline. After you update a random access volume to DESTROYED, you cannot vary the volume online.

If you update a sequential access volume to DESTROYED, the server does not attempt to mount the volume.

If a volume contains no files and you change the access mode to DESTROYED, the server deletes the volume from the database.

OFFSITE

Specifies that a copy, container-copy, or active-data storage pool volume is at an offsite location from which it cannot be mounted. Only volumes in copy, container-copy, or active-data storage pools can have the access mode of OFFSITE.

If you specify values for both the ACCESS and LOCATION parameters, but the access mode cannot be updated for a particular volume, the location attribute is also not updated for that volume. For example, if you specify ACCESS=OFFSITE and a LOCATION value for a primary storage pool volume, neither the access nor location values are updated because a primary storage pool volume cannot be given an access mode of OFFSITE.

LOCATION

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential access storage pools. The maximum length of the location is 255 characters. Enclose the location in quotation marks if it contains any blank characters. To remove a previously defined location, specify the null string ("").

WHERESTGPOOL

Specifies the name of the storage pool for volumes to be updated. Use this parameter to restrict the update by storage pool. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a storage pool name, volumes belonging to any storage pool are updated.

WHEREDEVCLASS

Specifies the name of the device class for volumes to be updated. Use this parameter to restrict the update by device class. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a device class name, volumes with any device class are updated.

WHEREACCESS

Specifies the current access mode of volumes to be updated. Use this parameter to restrict the update to volumes that currently have the specified access mode. This parameter is optional. You can specify multiple access modes by separating the modes with commas and no intervening spaces. If you do not specify a value for this parameter, the update is not restricted by the current access mode of a volume. Possible values are:

READWRITE

Update volumes with an access mode of READWRITE.

READONLY

Update volumes with an access mode of READONLY.

UNAVAILABLE

Update volumes with an access mode of UNAVAILABLE.

OFFSITE

Update volumes with an access mode of OFFSITE.

DESTROYED

Update volumes with an access mode of DESTROYED.

WHERESTATUS

Specifies the status of volumes to be updated. Use this parameter to restrict the update to volumes that have a specified status. This parameter is optional. You can specify multiple status values by separating the values with commas and no intervening spaces. If you do not specify a value for this parameter, the update is not restricted by volume status. Possible values are:

ONLINE

Update volumes with a status of ONLINE.

OFFLINE

- Update volumes with a status of OFFLINE.
- EMPTy Update volumes with a status of EMPTY.
- PENding Update volumes with a status of PENDING. These are volumes from which all files were deleted, but the time that is specified by the REUSEDELAY parameter has not elapsed.
- FILLing Update volumes with a status of FILLING.
- FULL Update volumes with a status of FULL.

Preview

Specifies whether you want to preview the update operation without updating volumes. This parameter is optional. The default value is NO. Possible values are:

- No Specifies that volumes are updated.
- Yes Specifies that you want only to preview the update operation. This option displays the volumes that will be updated if you run the update operation.

Example: Make a tape volume unavailable

Update a tape volume that is named DSMT20 to make it unavailable to client nodes and server processes.

```
update volume dsmt20 access=unavailable
```

Example: Update the access mode of all offsite volumes in a specific storage pool

Update all empty, offsite volumes in the TAPEPOOL2 storage pool. Set the access mode to READWRITE and delete the location information for the updated volumes.

```
update volume * access=readwrite location="" wherestgpool=tapepool2
  whereaccess=offsite wherestatus=empty
```

Related commands

Table 1. Commands related to UPDATE VOLUME

Command	Description
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
QUERY VOLUME	Displays information about storage pool volumes.
VARY	Specifies whether a disk volume is available to the server for use.

VALIDATE commands

Use the VALIDATE command to verify that an object is complete or valid for IBM Spectrum Protect™.

- **Linux** VALIDATE ASPERA (Validate an Aspera FASP configuration)
- **AIX Linux Windows** VALIDATE CLOUD (Validate cloud credentials)
- VALIDATE LANFREE (Validate LAN-Free paths)
- VALIDATE POLICYSET (Verify a policy set)
- VALIDATE REPLICATION (Validate replication for a client node)
- VALIDATE REPLPOLICY (Verify the policies on the target replication server)

Linux

VALIDATE ASPERA (Validate an Aspera FASP configuration)

Use this command to determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can be used to optimize data transfer in your system environment. Specifically, you can determine whether Aspera FASP technology would result in better network throughput than TCP/IP technology.

This command verifies the following additional items:

- Whether the system environment is correctly configured to use Aspera FASP technology
- Whether the required licenses for enabling Aspera FASP technology are installed

Aspera FASP technology is used to optimize data transfer for node replication or storage pool protection in a wide area network (WAN). However, you are not required to configure your system for node replication or storage pool protection to run the VALIDATE ASPERA command. If your system is configured for node replication or storage pool protection in a local environment, you can issue the command to evaluate whether the data can be successfully replicated to a remote server.

This command is available only on Linux x86_64 operating systems.

Before you issue the command, complete the following tasks:

1. Ensure that at least one server is defined in your system environment. Issue the PING SERVER command to ensure that you have connectivity to the defined server. For example, if the server is named VMRH6T, issue the following command:

```
ping server vmrh6t
```

2. To use the VALIDATE ASPERA command to determine the speed of network throughput, install 30-day evaluation licenses or full, unlimited licenses on the source and target servers. For example, install licenses on the source and target servers, VMRH6 and VMRH6T. For instructions about obtaining and installing licenses, see Determining whether Aspera FASP technology can optimize data transfer in your system environment.

To simulate an environment that uses multiple sessions, you can run several instances of the VALIDATE ASPERA command simultaneously. If you plan to run multiple sessions, you might want to limit the bandwidth of each network connection to ensure that sufficient bandwidth is available for all network connections. To limit the bandwidth, specify the FASPTARGETRATE server option as described in FASPTARGETRATE.

You can query the current transferred amount by issuing the QUERY PROCESS command:

```
query process
```

You can obtain the process number from the output of the QUERY PROCESS command. You can cancel the process by issuing the CANCEL PROCESS command and specifying the process number, for example:

```
cancel process 3
```

Privilege class

Any administrator can issue this command.

Syntax

```
>>-VALidate ASPera-+-----+----->
                    '---target_server_name---'
                    .-Wait---No-----
>--+-----+-----><
  '-DURation---seconds-' '-Wait---+No---+'
                          '-Yes-'
```

Parameters

target_server_name

Specifies a previously defined server. This parameter is optional. To specify this parameter, follow the guidelines:

- To determine whether Aspera FASP can optimize a node replication process, specify a target server that is configured for node replication.

- To determine whether Aspera FASP can optimize a storage pool protection process, specify a target server that is configured for storage pool protection.
- To determine whether Aspera FASP can optimize data transfer to a remote server that is defined but not configured for storage pool protection or node replication, specify that target server.
- If you do not specify a target server, the command output indicates whether the source server is correctly configured for Aspera FASP data transmission. The output also indicates whether a valid license for Aspera FASP is installed on the source server.

DURation

Specifies the allotted time, in seconds, for transferring data across the network to evaluate throughput. This parameter is optional. The default value is 120 seconds. You can specify a value in the range 120 - 3600000 seconds. The allotted time is divided between the Aspera FASP and TCPIP data transfers.

Wait

Specifies whether to wait for the server to complete the command processing. This parameter is optional. The default value is NO. You can specify one of the following values:

No

Specifies that the server processes the command in the background. You can continue with other tasks while the command is being processed. If you specify NO, the output messages are displayed in the activity log.

Yes

Specifies that the server processes the command in the foreground. The operation must complete processing before you can continue with other tasks. If you specify YES, the output messages are displayed in the administrative command-line client.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Display information about the status of an Aspera FASP configuration

On the source server, run the `VALIDATE ASPERA` command. To ensure that messages are displayed in the administrative command-line client, specify `WAIT=YES`. See Field descriptions for field descriptions.

```
validate aspera wait=yes
```

```
ANR3836I Validation of the Aspera FASP connection from VMRH6 to localhost.
Amount transferred using FASP: 0 MB per second. Amount transferred using
TCP/IP: 0 MB per second. Latency: 0 microseconds. Status: OK. Days until
license expires: Never.
```

Example: Verify whether the required licenses are installed

On the source server, run the `VALIDATE ASPERA` command and specify the target replication server. To ensure that messages are displayed in the administrative command-line client, specify `WAIT=YES`. See Field descriptions for field descriptions.

```
validate aspera vmrh6t wait=yes
```

```
ANR0984I Process 8 for VALIDATE ASPERA started in the FOREGROUND at 09:35:21 AM.
ANR3672E The license file that is required to enable Aspera Fast Adaptive
Secure Protocol (FASP) technology was not found on the VMRH6 server.
ANR3836I Validation of the Aspera FASP connection from VMRH6 to localhost.
Amount transferred using FASP: 0 MB per second. Amount transferred using
TCP/IP: 0 MB per second. Latency: 0 microseconds. Status: Invalid
configuration. Days until license expires: Expired.
ANR0985I Process 8 for VALIDATE ASPERA running in the FOREGROUND completed with
completion state FAILURE at 09:35:21 AM.
ANR1893E Process 8 for VALIDATE ASPERA completed with a completion state of
FAILURE.
```

Field descriptions

Status

The status of the configuration. The following values are possible:

- `OK` indicates that no issues are detected.
- `Invalid configuration` indicates that a configuration file, license file, or Aspera FASP library file is missing.
- `License issue` indicates that a license is missing, invalid, or expired.

- `Server failure` indicates that all ports are in use, a network read/write error occurred, or the Aspera FASP log file is unwritable.
- `Invalid target configuration` indicates that a configuration file, license file, or Aspera FASP library file is missing on the target server.
- `Failure on target server` indicates that all ports are in use, a network read/write error occurred, or the Aspera FASP log file is unwritable.
- `License issue on target server` indicates that a license is invalid or expired on the target server.
- `Unsupported operating system` indicates that an operating system other than Linux x86_64 is installed on one or both servers.
- `Unknown` indicates that an unexpected error occurred. To identify the error, review the log messages.

Days until license expires

The following values are possible:

- `Never` indicates that a full, unlimited license is installed.
- `Today` indicates that a 30-day evaluation license is installed and it expires today.
- `Expired` indicates that a 30-day evaluation license is installed, but has expired.
- `Number` indicates that a 30-day evaluation license is installed and will expire in the specified number of days.
- `License not found` indicates that no license was found.

Amount transferred using TCP/IP

The speed of data transfer, in megabytes per second, using TCP/IP technology.

Amount transferred using FASP

The speed of data transfer, in megabytes per second, using Aspera FASP technology.

Latency

The latency of data transfer in microseconds.

Related commands

Table 1. Commands related to VALIDATE ASPERA

Command	Description						
CANCEL SESSION	Cancels active sessions with the server.						
DEFINE SERVER	Defines a server for server-to-server communications.						
PING SERVER	Tests the connections between servers.						
<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>AIX</td><td>Linux</td><td>Windows</td></tr></table> PROTECT STGPOOL	AIX	Linux	Windows	<table border="1" style="display: inline-table; vertical-align: middle;"><tr><td>AIX</td><td>Linux</td><td>Windows</td></tr></table> Protects a directory-container storage pool.	AIX	Linux	Windows
AIX	Linux	Windows					
AIX	Linux	Windows					
REPLICATE NODE	Replicates data in file spaces that belong to a client node.						

VALIDATE CLOUD (Validate cloud credentials)

Before you define a storage pool, use this command to ensure that the credentials for a cloud-container storage pool are valid and that the necessary permissions are granted to the user.

Privilege class

Any administrator can issue this command.

Syntax

```

      .-CLOUDType---Swift-----
>>-VALIDATE CLOUD--+-+-----+----->
      '-CLOUDType---+Azure-----+'
                        +-S3-----+
                        +-IBMCloudswift-+
                        +-Swift-----+
                        '-V1Swift-----'

                                     (1)
>--CLOUDUrl---cloud_url--IDentity---cloud_identity----->

```

```
>--PAssword---password-----+-----+-----><
|                                     (2) |
'--BUCKETName---bucket_name-----'
```

Notes:

1. If you specify CLOUDTYPE=AZURE, do not specify the IDENTITY parameter.
2. The BUCKETNAME parameter is valid only if you specify CLOUDTYPE=S3.

Parameters

CLOUDType

Specifies the type of cloud environment where you are configuring the storage pool.
You can specify one of the following values:

AZure

Specifies that the storage pool uses a Microsoft Azure cloud computing system.

S3

Specifies that the storage pool uses a cloud computing system with the Simple Storage Service (S3) protocol, such as IBM® Cloud Object Storage or Amazon Web Services (AWS) S3.

IBMCloudswift

Specifies that the storage pool uses an IBM Cloud cloud computing system with an OpenStack Swift cloud computing system.

Swift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 2 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol that it is using.

V1Swift

Specifies that the storage pool uses an OpenStack Swift cloud computing system. This value also specifies that the storage pool uses Version 1 of the protocol for authentication to the cloud. The URL of the cloud usually contains the version number of the protocol that it is using.

This parameter is optional. If you do not specify the parameter, the default value, SWIFT, is used.

CLOUDUrl (Required)

Specifies the URL of the cloud environment where you configure the storage pool. Based on your cloud provider, you can use a blob service endpoint, region endpoint URL, an accesser IP address, a public authentication endpoint, or a similar value for this parameter. Ensure that you include the protocol, such as `https://` or `http://`, at the beginning of the URL. The maximum length of the web address is 870 characters. The CLOUDURL parameter is validated when the first backup begins.

IDentity (Required)

Specifies the user ID for the cloud. This parameter is required for all supported cloud computing systems except Azure. If you specify CLOUDTYPE=AZURE, do not specify the IDENTITY parameter. Based on your cloud provider, you can use an access key ID, a user name, a tenant name and user name, or a similar value for this parameter. The maximum length of the user ID is 255 characters.

PAssword (Required)

Specifies the password for the cloud. Based on your cloud provider, you can use a shared access signature (SAS) token, secret access key, an API key, a password, or a similar value for this parameter. This parameter is required. The maximum length of the password is 255 characters.

BUCKETName

Specifies the name for an AWS S3 bucket or a IBM Cloud Object Storage vault to use with this storage pool, instead of using the default bucket name or vault name. This parameter is optional, and is valid only if you specify CLOUDTYPE=S3. If a bucket or vault exists with the name that you specify, that bucket or vault is tested to ensure that the proper permissions are set. If the bucket or vault does not exist, the parameter verifies only that a bucket or vault with that name does not exist. Follow the naming restrictions for your cloud provider when you specify this parameter. Review the permissions for the bucket or vault and make sure that the credentials have permission to read, write, list, and delete objects in this bucket or vault.

Tip: If you do not specify the BUCKETNAME parameter, the Replication Globally Unique ID is used as the default bucket name. The default is

`ibmsp guid`

where *guid* is the REPLICATION GLOBALLY UNIQUE ID value, minus the periods, in the output of the QUERY REPLSERVER command. For example, if the Replication Globally Unique ID is 52.82.39.20.64.d0.11.e6.9d.77.0a.00.27.00.00.00, the default bucket name is `ibmsp.5282392064d011e69d770a0027000000`.

Example: Verify the credentials of an S3 cloud-container storage pool

Validate the credentials of the cloud-container storage pool.

```
validate cloud
cloudtype=s3 clouurl=http://123.234.123.234:5000/v2.0
password=protect8991 bucketname=ibmsp.5282392064d011e69d770a0027000000
```

Related commands

Table 1. Commands related to VALIDATE CLOUD

Command	Description
DEFINE STGPOOL (cloud-container)	Define a cloud-container storage pool.
QUERY REPLSERVER	Displays information about replicating servers.
UPDATE STGPOOL (cloud-container)	Update a cloud-container storage pool.

VALIDATE LANFREE (Validate LAN-Free paths)

Use this command to determine which destinations for a given node using a specific storage agent are capable of LAN-Free data movement.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
>>-VALidate LANfree--node_name--stgagent_name-----<<
```

Parameters

`node_name` (Required)

The name of the node to evaluate.

`stgagent_name` (Required)

The name of the storage agent to evaluate.

Example: Validate a current LAN-Free configuration

Validate the current server definitions and configuration for node TIGER to use storage agent AIX_STA1 for LAN-free data operations.

```
validate lanfree tiger aix_sta1
```

Node Name	Storage Agent	Operation	Mgmt Class Name	Destination Name	LAN-Free capable?	Explanation
TIGER	AIX_STA1	BACKUP	STANDARD	OUTPOOL	NO	No available online paths.
TIGER	AIX_STA1	BACKUP	STANDARD	PRIMARY	NO	Destination storage pool is configured for simultaneous write.
TIGER	AIX_STA1	BACKUP	STANDARD	SHRPOOL	YES	
TIGER	AIX_STA1	BACKUP	NOARCH	LFFILE	NO	Storage pool contains data

TIGER AIX_STA1 ARCHIVE	STANDARD	OUTPOOL	NO	deduplicated by clients, and is not accessible by storage agents V6.1 or earlier. No available online paths. Destination storage pool is configured for simultaneous write.
TIGER AIX_STA1 ARCHIVE	STANDARD	PRIMARY	NO	
TIGER AIX_STA1 ARCHIVE	STANDARD	SHRPOOL	YES	

Related commands

Table 1. Commands related to VALIDATE LANFREE

Command	Description
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY DEVCLASS	Displays information about device classes.
QUERY DOMAIN	Displays information about policy domains.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY MGMTCLASS	Displays information about management classes.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PATH	Displays information about the path from a source to a destination.
QUERY POLICYSET	Displays information about policy sets.
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
QUERY STGPOOL	Displays information about storage pools.

VALIDATE POLICYSET (Verify a policy set)

Use this command to verify that a policy set is complete and valid before you activate it. The command examines the management class and copy group definitions in the policy set and reports on conditions that you need to consider before activating the policy set.

The VALIDATE POLICYSET command fails if any of the following conditions exist:

- The policy set has no default management class.
- A copy group within the policy set specifies a copy storage pool as a destination.
- A management class specifies a copy storage pool as the destination for files that were migrated by an IBM Spectrum Protect™ for Space Management client.
- A TOCDESTINATION parameter is specified, and the storage pool is either a copy pool or has a data format other than NATIVE or NONBLOCK.

The server issues warning messages for the following conditions:

- A copy group specifies a storage pool that does not exist as a destination for backed-up or archived files.

If you activate a policy set with copy groups that specify nonexistent storage pools, the client backup or archive operations fail.

- A management class specifies a storage pool that does not exist as a destination for files migrated by IBM Spectrum Protect for Space Management clients.
- The policy set does not have one or more management classes that exist in the current ACTIVE policy set.

If you activate the policy set, backup files bound to the deleted management classes are rebound to the default management class in the new active policy set.

- The policy set does not have one or more copy groups that exist in the current ACTIVE policy set.

If you activate the policy set, files bound to the management classes with deleted copy groups are no longer archived or backed up.

- The default management class for the policy set does not contain a backup or archive copy group.

If you activate the policy set with this default management class, clients using the default cannot back up or archive files.

- A management class specifies that a backup version must exist before a file can be migrated from a client node (MIGREQUIRESBKUP=YES), but the management class does not contain a backup copy group.

If the server has data retention protection enabled, the following conditions must exist:

- All management classes in the policy set to be validated must contain an archive copy group.
- If a management class exists in the active policy set, a management class with the same name must exist in the policy set to be validated.
- If an archive copy group exists in the active policy set, the corresponding copy group in the policy set to be validated must have a RETVER value at least as large as the corresponding values in the active copy group.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
>>-VALIDATE Policyset--domain_name--policy_set_name-----><
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the policy set is assigned.

policy_set_name (Required)

Specifies the name of the policy set to be validated.

Example: Validate a specific policy set

Validate the policy set VACATION located in the EMPLOYEE_RECORDS policy domain.

```
validate policyset employee_records vacation
```

Related commands

Table 1. Commands related to VALIDATE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.

Command	Description
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE POLICYSET	Changes the description of a policy set.

VALIDATE REPLICATION (Validate replication for a client node)

Use this command to identify the replication rules that apply to file spaces in client nodes that are configured for replication. You can also use this command to verify that the source replication server can communicate with the target replication server.

Before you begin replication processing, use the VALIDATE REPLICATION command to determine whether your replication configuration is correct.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

      .-,------.
      v          |
>>-VALIDate REPLication-----node_name---+----->
      .-VERIFYconnection-----No-----
>--+-----+-----+-----+----->>
      '-VERIFYconnection-----+No--+-'
      '-Yes-'

```

Parameters

node_name (Required)

Specifies the name of the client node whose file spaces you want to display. To specify multiple client node names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify names.

Information is displayed only for client nodes that are either enabled or disabled for replication. The replication mode must be SEND. To determine whether a client node is enabled or disabled for replication and its mode, issue the QUERY NODE command. Look for values in the Replication State and Replication Mode fields.

VERIFYconnection

Specifies whether to check the connection to a target replication server. The version of the target replication server is also checked to verify that it is Version 6.3 or later. This parameter is optional. The default is NO. You can specify one of the following values:

No

The connection and version of the target replication server are not checked.

Yes

The connection and version of the target replication server are checked.

Example: Validate replication for a client node

The name of the client node is NODE1. Verify the connection status between the source and the target replication servers.

```
validate replication node1 verifyconnection=yes
```

```

Node Name: NODE1
Filespace Name: \\node1\c$

```

```

        FSID: 1
        Type: Bkup
Controlling Replication Rule: ACTIVE_DATA
    Replication Rule Level: System Level
        Server Name: DRSRV
    Connection Status: Valid Connection

        Node Name: NODE1
    Filespace Name: \\node1\c$
        FSID: 1
        Type: Arch
Controlling Replication Rule: ALL_DATA_HIGH_PRIORITY
    Replication Rule Level: Node Level
        Server Name: DRSRV
    Connection Status: Valid Connection

        Node Name: NODE1
    Filespace Name: \\node1\c$
        FSID: 1
        Type: SpMg
Controlling Replication Rule: ALL_DATA
    Replication Rule Level: System Level
        Server Name: DRSRV
    Connection Status: Valid Connection

```

Output is displayed for all data types regardless of whether a file space contains the data types. For example, if a file space contains only backup and archive data, the output of the VALIDATE REPLICATION command also contains information that would be relevant to space-managed data.

Field descriptions

Node Name

The node that owns the replicated data.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Operations Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, the name is converted to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

FSID

The file space identifier for the file space. The server assigns a unique FSID when a file space is first stored on the server.

Type

The type of data. The following values are possible:

Arch

Archive data

Bkup

Backup data

SpMg

Data that was migrated by an IBM Spectrum Protect™ for Space Management client.

Controlling Replication Rule

The name of the replication rule that controls replication for a data type in a file space. To determine whether the controlling rule is a file space rule, a client rule, or a server rule, check the Replication Rule Level field.

Replication Rule Level

The level of the controlling rule in the replication-rule hierarchy. The following values are possible:

Filespace

The controlling rule is assigned to a data type in the file space.

Node

The controlling rule is assigned to a data type for a client node.

Server

The controlling rule is assigned to a data type for all file spaces in all client nodes that are configured for replication.

Server Name

The name of the target replication server to be queried.

Connection Status

The connection status between the source and the target replication server. The following values are possible:

Valid Connection

Communication with the target replication server was successful, and the target replication server is a V6.3 server.

Target Server Not Set

The target replication server is not set. To set the target replication server, issue the SET REPLSERVER command.

Communication Failure

The source replication server was unable to contact the target replication server. Examine the activity log for error messages about failed communications. Consider the following possible causes:

- The replication configuration on the source replication server is not valid. One or more of the following problems might exist:
 - The server definition for the target replication server is incorrect.
 - If the target replication-server definition was deleted and redefined, issue the PING SERVER command to test the connection between the source and the target replication server. If the PING SERVER command is successful, issue the UPDATE SERVER command and specify FORCESYNC=YES to reset the server verification keys.
 - The server name, server low-level address, server high-level address, and server password do not match the values that are specified in the server definition on the target replication server.
- The replication configuration on the target replication server is not valid. One or more of the following problems might exist:
 - The version of the target replication server is earlier than V6.3.
 - The server definition for the source replication server is incorrect.
 - The server name, server low-level address, server high-level address, and server password do not match the values that are specified in the server definition on the source replication server.
- Network communications are unavailable. To test the connection between the source and target server, issue the PING SERVER command.
- The target replication server is unavailable.
- Sessions between the source and the target replication servers are disabled. To verify the status of sessions, issue the QUERY STATUS command.

Replication Suspended

Replication processing is suspended when you restore the database on the source replication server or you disable replication processing on this server by issuing the DISABLE REPLICATION command.

Related commands

Table 1. Commands related to VALIDATE REPLICATION

Command	Description
DISABLE REPLICATION	Prevents outbound replication processing on a server.
ENABLE REPLICATION	Allows outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLRULE	Displays information about node replication rules.
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

Command	Description
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
SET REPLSERVER	Specifies a target replication server.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes that are associated with a client node.
UPDATE REPLRULE	Enables or disables replication rules.
UPDATE SERVER	Updates information about a server.

VALIDATE REPLPOLICY (Verify the policies on the target replication server)

Use this command to compare the policies for client nodes on the source replication server with the same policies on the target replication server where the client node data is being replicated.

The command displays the differences between these policies so that you can verify that any differences between the policies on the source and target replication servers are intended or you can modify the policies on the target replication server.

Ensure that IBM Spectrum Protect™, Version 7.1.1 or later, is installed on the source and target replication servers before you issue this command. Issue this command on the source replication server.

Privilege class

Any administrator can issue this command.

Syntax

```
>>-VALIDate REPLPolicy--+-+-----+----->>
                        '-server_name-'
```

Parameters

server_name

Specifies the name of the target replication server that has policies you want to verify. This parameter is optional. If you do not specify this parameter, the command sets the default replication server as the target replication server.

Example: Display the differences between the replication policies on a source and target replication server

To display the differences between the policies on the source replication server and the policies on the target replication server, CVTCVS_LXS_SRV2, where the client data is replicated, issue the following command on the source replication server:

```
VALIDATE REPLPOLICY CVTCVS_LXS_SRV2
```

Policy domain name on this server	Policy domain name on target server	Target server name
-----	-----	-----
STANDARD	STANDARD	CVTCVS_LXS_SRV2
Differences in policy set:		
Change detected	Source server value	Target server value
-----	-----	-----
Mgmt class only on target	Not applicable	STANDARD2

Mgmt Class only on source	STANDARD1	Not applicable
Differences in backup copy group	STANDARD in management class	STANDARD
Change detected	Source server value	Target server value

Versions data exists	2	20
Affected nodes		

NODE1,NODE2,NODE3,NODE4,NODE5		

Field descriptions

Policy domain name on this server

Specifies the policy domain name on the source replication server where the command is issued.

Policy domain name on target server

Specifies the policy domain name on the target replication server.

Target server name

Specifies the name of the target replication server.

Differences in policy set:

Specifies the differences between the policies that are defined on the source and target replication servers. The differences between the policies are listed under the following fields:

Change detected

Specifies the list of policy items that are different between the source and target replication servers.

Source server value

Specifies the value for the policy item on the source replication server.

Target server value

Specifies the value for the policy item on the target replication server.

Differences in backup copy group <backup_copy_group_name> in default management class OR Differences in archive copy group <archive_copy_group_name> in default management class

Specifies the differences between the backup copy group or the archive copy group in the management class. The differences are listed under the following fields:

Change Detected

Specifies the list of copy group fields that are different.

Source server value

Specifies the value in the copy group field on the source replication server.

Target server value

Specifies the value in the copy group field on the target replication server.

Affected nodes

Specifies the names of all the client nodes that are affected by the changes that are shown in this output.

Related commands

Table 1. Commands related to VALIDATE REPLPOLICY

Command	Description
VALIDATE REPLICATION	Verifies replication for file spaces and data types.
QUERY REPLSERVER	Displays information about replicating servers.
SET DISSIMILARPOLICIES	Enable the policies on the target replication server to manage replicated data.
QUERY DOMAIN	Displays information about policy domains.
QUERY POLICYSET	Displays information about policy sets.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.

VARY (Bring a random access volume online or offline)

Use this command to make a random access storage pool volume online or offline to the server.

Privilege class

This command is valid only for volumes on random access devices. For example, use this command during maintenance or corrective action of a random access volume. You cannot vary a random access volume online that is defined as unavailable.

To issue this command, you must have system privilege or operator privilege.

Syntax

```
>>-VArY--+-ONline--+-+volume_name--+-+-----+-----><
      '-Offline-'      '-Wait-----+No--+-'
                          '-Yes-'
```

Parameters

ONline

Specifies that the server can use the random access volume.

OFFline

Specifies that the server cannot use the volume.

volume_name (Required)

Specifies the volume identifier. Volume names cannot contain embedded blanks or equal signs.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No

Specifies that the server processes this command in the background, while other tasks run. The server displays messages created from the background process either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before you continue with other tasks. The server displays the output messages to the administrative client when the command completes.

AIX | **Linux** | **Windows** You cannot specify WAIT=YES from the server console.

Example: Bring volume online

AIX | **Linux** Make volume /adsm/stgvol/1 available to the server for use as a storage pool volume. **AIX** | **Linux**

```
vary online /adsm/stgvol/1
```

Windows Make volume j:\storage\pool001 available to the server for use as a storage pool volume. **Windows**

```
vary online j:\storage\pool001
```

Related commands

Table 1. Commands related to VARY

Command	Description
CANCEL PROCESS	Cancel a background server process.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.

Command	Description
DELETE VOLUME	Deletes a volume from a storage pool.
QUERY PROCESS	Displays information about background processes.
QUERY VOLUME	Displays information about storage pool volumes.

Server options

At installation, IBM Spectrum Protect™ provides a server options file that contains a set of default options to start the server.

The file is:

- dsmserv.opt in the server instance directory

Server options let you customize the following:

- Communication
- Server storage
- Client-server
- Date, number, time, and language
- Database and recovery log
- Data transfer
- Message
- Event logging
- Security and licensing

Several other options are available for miscellaneous purposes. These undocumented options are intended to be used only by IBM® support.

To display the current option settings, enter:

```
query option
```

- Modifying server options
The server reads the server options file at server initialization. When you update a server option by editing the file, you must stop and start the server to activate the updated server options file.
- Types of server options
Server options let you customize how some functions and processes work.
- 3494SHARED
The 3494SHARED option specifies whether an IBM 3494 library can share applications other than IBM Spectrum Protect.
- ACSACCESSID
The ACSACCESSID option specifies the ID for the ACS access control for an ACSLS library.
- ACSLOCKDRIVE
The ACSLOCKDRIVE option specifies if the drives within the ACSLS libraries are locked. Drive locking ensures the exclusive use of the drive in the ACSLS library in a shared environment. However, there is some performance gain if libraries are not locked. When other applications do not share the IBM Spectrum Protect drives, drive locking is not required.
- ACSQUICKINIT
The ACSQUICKINIT option specifies whether, at server startup, the initialization of the ACSLS library is a quick or full initialization. The default is Yes. A quick initialization avoids the overhead associated with synchronizing the IBM Spectrum Protect server inventory with the ACSLS library inventory (through an audit of the library).
- ACSTIMEOUTX
The ACSTIMEOUTX option specifies the multiple for the built-in timeout value for ACSLS APIs. The built-in timeout value for the ENTER, EJECT, and AUDIT ACS API is 1800 seconds; for all other ACSLS APIs it is 600 seconds. For example, if the multiple value specified is 5, the timeout value for audit API becomes 9000 seconds, and all other APIs become 3000 seconds.
- ACTIVELOGDIRECTORY
The ACTIVELOGDIRECTORY option specifies the name of the directory where all active logs are stored.
- ACTIVELOGSIZE
The ACTIVELOGSIZE option sets the total log size.
- ADMINCOMMTIMEOUT
The ADMINCOMMTIMEOUT option specifies how long the server waits for an expected administrative client message during an operation that causes a database update.

- **ADMINIDLETIMEOUT**
The ADMINIDLETIMEOUT option specifies the amount of time, in minutes, that an administrative client session can be idle before the server cancels the session.
- **ADMINONCLIENTPORT**
The ADMINONCLIENTPORT option specifies whether the TCPSPORT can be used by administrative sessions. The default is YES.
- **Windows ADSMGROUPNAME**
The ADSMGROUPNAME option specifies the name of a Windows group. A client node must be a member of this group to use the IBM Spectrum Protect server through NT Unified Logon. The client node must also be a registered IBM Spectrum Protect client node.
- **ALIASHALT**
The ALIASHALT option allows administrators to give the IBM Spectrum Protect **HALT** command a different name.
- **ALLOWDESAUTH**
The ALLOWDESAUTH option specifies whether to allow use of the Data Encryption Standard (DES) algorithm for authentication between a server and a backup-archive client.
- **ALLOWREORGINDEX**
The ALLOWREORGINDEX option specifies whether server-initiated index reorganization is enabled or disabled.
- **ALLOWREORGTABLE**
The ALLOWREORGTABLE option specifies whether server-initiated table reorganization is enabled or disabled.
- **ARCHFAILOVERLOGDIRECTORY**
The ARCHFAILOVERLOGDIRECTORY option specifies the directory which the server uses to store archive log files that cannot be stored in the archive log directory.
- **ARCHLOGCOMPRESS**
You can enable or disable compression of archive logs on the IBM Spectrum Protect server. By compressing the archive logs, you reduce the amount of space that is required for storage.
- **ARCHLOGDIRECTORY**
The ARCHLOGDIRECTORY option specifies a directory that the database manager can archive a log file into after all the transactions represented in that log file are completed.
- **ARCHLOGUSEDTHRESHOLD**
The ARCHLOGUSEDTHRESHOLD option specifies when to start an automatic database backup in relation to the percentage of archive log file space used. The default is 80 percent.
- **ASSISTVCRRECOVERY**
The ASSISTVCRRECOVERY option specifies whether IBM Spectrum Protect assists an IBM 3590 drive in recovering from a lost or corrupted Vital Cartridge Records (VCR) condition. If you specify YES (the default) and if IBM Spectrum Protect detects an error during the mount processing, it locates to the end-of-data during the dismount processing to allow the drives to restore the VCR. During the tape operation, there might be some small effect on performance because the drive cannot complete a fast locate with a lost or corrupted VCR. However, there is no loss of data.
- **AUDITSTORAGE**
As part of a license audit operation, the server calculates, by node, the amount of server storage used for backup, archive, and space-managed files. For servers managing large amounts of data, this calculation can take a great deal of CPU time and can stall other server activity. You can use the AUDITSTORAGE option to specify that storage is not to be calculated as part of a license audit.
- **BACKUPINITIATIONROOT**
The BACKUPINITIATIONROOT option specifies whether the server overrides node parameter values for users who are not IBM Spectrum Protect authorized users.
- **CHECKTAPEPOS**
The CHECKTAPEPOS option specifies whether the IBM Spectrum Protect server validates the position of data blocks on tape.
- **CLIENTDEDUPTXNLIMIT**
The CLIENTDEDUPTXNLIMIT option specifies the maximum size of a transaction when client-side deduplicated data is backed up or archived.
- **CLIENTDEPLOYCATALOGURL**
The CLIENTDEPLOYCATALOGURL option specifies the location of the catalog file that is used for automatic client deployment operations.
- **CLIENTDEPLOYUSELOCALCATALOG**
The CLIENTDEPLOYCATALOGURL option specifies whether the local version of the catalog file is used for automatic client deployment operations.
- **COMMMETHOD**
The COMMMETHOD option specifies a communication method to be used by the server.
- **COMMTIMEOUT**
The COMMTIMEOUT option specifies how long the server waits for an expected client message during an operation that causes a database update. If the length of time exceeds this time-out, the server ends the session with the client. You may

want to increase the time-out value to prevent clients from timing out. Clients may time out if there is a heavy network load in your environment or they are backing up large files.

- **CONTAINERRESOURCE_TIMEOUT**
The CONTAINERRESOURCE_TIMEOUT option specifies how long the server waits to complete a data store operation to a container storage pool.
- **Windows DATEFORMAT**
The DATEFORMAT option specifies the format in which dates are displayed by the server.
- **DBDIAGLOGSIZE**
This option helps to control the amount of space that is used by diagnostic log files.
- **DBDIAGPATHFSTHRESHOLD**
The DBDIAGPATHFSTHRESHOLD option specifies the threshold for free space on the file system or disk that contains the db2diag.log file.
- **DBMEMPERCENT**
Use this option to specify the percentage of the virtual address space that is dedicated to the database manager processes.
- **DBMTCPPORT**
The DBMTCPPORT option specifies the port number on which the TCP/IP communication driver for the database manager waits for requests for client sessions.
- **DEDUPREQUIRESBACKUP**
The DEDUPREQUIRESBACKUP option specifies whether volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and whether duplicate data can be discarded before the storage pools are backed up.
- **DEDUPTIER2FILESIZE**
The DEDUPTIER2FILESIZE option specifies at what file size IBM Spectrum Protect begins to use Tier 2 data deduplication.
- **DEDUPTIER3FILESIZE**
The DEDUPTIER3FILESIZE option specifies at what file size IBM Spectrum Protect begins to use Tier 3 data deduplication.
- **DEVCONFIG**
The DEVCONFIG option specifies the name of a file in which you want IBM Spectrum Protect to store a backup copy of device configuration information.
- **DISABLEREORGTABLE**
The DISABLEREORGTABLE option specifies whether online table reorganization is disabled for table names that are specified in the tables list.
- **DISABLESCHEDS**
The DISABLESCHEDS option specifies whether administrative and client schedules are disabled during IBM Spectrum Protect server recovery.
- **DISPLAYLFINFO**
The DISPLAYLFINFO option specifies how the accounting records and summary table entries report the node name.
- **DNSLOOKUP**
The DNSLOOKUP option specifies whether the server uses system API calls to determine the domain name server (DNS) names of systems that contact the server.
- **DRIVEACQUIRERETRY**
The DRIVEACQUIRERETRY option lets you specify how many times the server retries the acquisition of a drive in an IBM 349x library. If the library is shared among multiple applications, its drives may appear to be available to the server (through the use of a background polling process) when they are not.
- **ENABLENASDEDUP**
The ENABLENASDEDUP server option specifies whether the server deduplicates data that is stored by a network-attached storage (NAS) file server. This option applies only to NetApp file servers.
- **EVENTSERVER**
The EVENTSERVER option specifies whether at startup the server should try to contact the event server.
- **EXPINTERVAL**
The EXPINTERVAL option specifies the interval, in hours, between automatic inventory expiration processes by IBM Spectrum Protect. Inventory expiration removes client backup and archive file copies from the server as specified by the management classes to which the client files are bound. If expiration is not run periodically, storage pool space is not reclaimed from expired client files, and the server requires more storage space than required by policy.
- **EXPQUIET**
The EXPQUIET option specifies whether IBM Spectrum Protect sends detailed messages during expiration processing.
- **Linux Windows FASPBEGPORT**
The FASPBEGPORT option specifies the starting number in the range of port numbers that are used for network communications with Aspera® Fast Adaptive Secure Protocol (FASP®) technology.
- **Linux Windows FASPENDPORT**
The FASPENDPORT option specifies the ending number in the range of port numbers that are used for network communications with Aspera Fast Adaptive Secure Protocol (FASP) technology.

- **Linux | Windows FASPTARGETRATE**
The FASPTARGETRATE option specifies the target rate for data transfer with Aspera Fast Adaptive Secure Protocol (FASP) technology. By specifying the target rate, you limit the bandwidth of each network connection that uses Aspera FASP technology. In this way, you can ensure that sufficient bandwidth is available for all network connections.
- **FFDCLOGLEVEL**
The FFDCLOGLEVEL option specifies the type of general server messages that are displayed in the first failure data capture (FFDC) log.
- **FFDCLOGNAME**
The FFDCLOGNAME option specifies a name for the first failure data capture (FFDC) log.
- **FFDCMAXLOGSIZE**
The FFDCMAXLOGSIZE option specifies the size for the first failure data capture (FFDC) log file.
- **FFDCNUMLOGS**
The FFDCNUMLOGS option specifies the number of log files that can be used for circular logging. The default value is 10.
- **FILEEXIT**
The FILEEXIT option specifies a file to which enabled events are routed. Each logged event is a record in the file.
- **FILETEXTEXIT**
The FILETEXTEXIT option specifies a file to which enabled events are routed. Each logged event is a fixed-size, readable line.
- **FIPSMODE**
The FIPSMODE option specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for non-Secure Sockets Layer (SSL) operations.
- **FSUSEDTHRESHOLD**
The FSUSEDTHRESHOLD option specifies what percentage of the file system can be filled up by the database before an alert message is issued.
- **IDLETIMEOUT**
The IDLETIMEOUT option specifies the amount of time, in minutes, that a client session can be idle before the server cancels the session. You may want to increase the time-out value to prevent clients from timing out if there is a heavy network load in your environment. Note, however, that a large number of idle sessions could prevent other users from connecting to the server.
- **KEEPALIVE**
The KEEPALIVE option specifies whether the Transmission Control Protocol (TCP) keepalive function is enabled for outbound TCP sockets. The TCP keepalive function sends a transmission from one device to another to check that the link between the two devices is operating.
- **KEEPALIVETIME**
The KEEPALIVETIME option specifies how often TCP sends a keepalive transmission when it receives a response. This option applies only if you set the KEEPALIVE option to YES.
- **KEEPALIVEINTERVAL**
The KEEPALIVEINTERVAL option specifies how often a keepalive transmission is sent if no response is received. This option applies only if you set the KEEPALIVE option to YES.
- **LANGUAGE**
The LANGUAGE option controls the initialization of locales. A locale includes the language and the date, time, and number formats to be used for the console and server.
- **LDAPCACHEDURATION**
The LDAPCACHEDURATION option determines the amount of time that the IBM Spectrum Protect server caches LDAP password authentication information.
- **LDAPURL**
The LDAPURL option specifies the location of a Lightweight Directory Access Protocol (LDAP) server. Set the LDAPURL option after you configure the LDAP server.
- **MAXSESSIONS**
The MAXSESSIONS option specifies the maximum number of simultaneous client sessions that can connect with the server.
- **MESSAGEFORMAT**
The MESSAGEFORMAT option specifies whether a message number is displayed in all lines of a multi-line message.
- **MIRRORLOGDIRECTORY**
The MIRRORLOGDIRECTORY option specifies the directory for mirroring the active log path.
- **MOVEBATCHSIZE**
The MOVEBATCHSIZE option specifies the number of client files that are to be moved and grouped together in a batch, within the same server transaction. This data movement results from storage pool backups and restores, migration, reclamation, and MOVE DATA operations. This option works with the MOVESIZETHRESH option.
- **MOVESIZETHRESH**
The MOVESIZETHRESH option specifies, in megabytes, a threshold for the amount of data moved as a batch, within the

same server transaction. When this threshold is reached, no more files are added to the current batch, and a new transaction is started after the current batch is moved.

- **MSGINTERVAL**
The MSGINTERVAL option specifies the time, in minutes, between messages prompting an operator to mount a tape for the server.
- **Windows** **NAMEDPIPENAME**
The NAMEDPIPENAME option specifies a communication method that allows processes to communicate with one another without having to know where the sender and receiver processes are located. The name acts like an alias, connecting the two processes regardless of whether they are on the same computer or across connected domains.
- **NDMPCONNECTIONTIMEOUT**
The NDMPCONNECTIONTIMEOUT server option specifies the time in hours that IBM Spectrum Protect server waits to receive status updates during NDMP restore operations across the LAN. NDMP restore operations of large NAS file systems can have long periods of inactivity. The default is 6 hours.
- **NDMPCONTROLPORT**
The NDMPCONTROLPORT option specifies the port number to be used for internal communications for certain Network Data Management Protocol (NDMP) operations. The IBM Spectrum Protect server does not function as a general purpose NDMP tape server.
- **NDMPENABLEKEEPALIVE**
The NDMPENABLEKEEPALIVE server option specifies whether the IBM Spectrum Protect server enables Transmission Control Protocol (TCP) keepalive on network data-management protocol (NDMP) control connections to network-attached storage (NAS) devices. The default is NO.
- **AIX** **Linux** **Windows** **NDMPKEEPIDLEMINUTES**
The NDMPKEEPIDLEMINUTES server option specifies the amount of time, in minutes, before the operating system transmits the first Transmission Control Protocol (TCP) keepalive packet on a network data-management protocol (NDMP) control connection. The default is 120 minutes.
- **NDMPPORTRANGE**
The NDMPPORTRANGE option specifies the range of port numbers through which IBM Spectrum Protect cycles to obtain a port number for accepting a session from a network-attached storage (NAS) device for data transfer. The default is 0,0 which means that IBM Spectrum Protect lets the operating system provide a port (ephemeral port).
- **NDMPPREFDATAINTERFACE**
This option specifies the IP address that is associated with the interface in which you want the server to receive all Network Data Management Protocol (NDMP) backup data.
- **NOPREEMPT**
The server allows certain operations to preempt other operations for access to volumes and devices. You can specify the NOPREEMPT option to disable preemption. When preemption is disabled, no operation can preempt another for access to a volume, and only a database backup operation can preempt another operation for access to a device.
- **NORETRIEVEDATE**
The NORETRIEVEDATE option specifies that the server does not update the retrieve date of a file in a disk storage pool when a client restores or retrieves the file. This option and the MIGDELAY storage pool parameter control when the server migrates files.
- **Windows** **NPAUDITFAILURE**
The NPAUDITFAILURE option specifies whether an event is sent to the event log when a node logs in to the server using a name that is in the Windows group but does not match the Windows account login name. To ensure that a node can access only its own data, the node name and the Windows account name must match.
- **Windows** **NPAUDITSUCCESS**
The NPAUDITSUCCESS option specifies that an event is sent to the event log when a client node user is authenticated for access to the server through SECUREPIPE.
- **Windows** **NPBUFFERSIZE**
The NPBUFFERSIZE option specifies the size of the Named Pipes communication buffer.
- **Windows** **NUMBERFORMAT**
The NUMBERFORMAT option specifies the format in which the server displays numbers.
- **NUMOPENVOLSALLOWED**
The NUMOPENVOLSALLOWED option specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time.
- **PUSHSTATUS**
The PUSHSTATUS option is used on spoke servers to ensure that status information is sent to the hub server. Do not update this option unless you must restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect servers are not defined as hub or spoke servers.
- **QUERYAUTH**
The QUERYAUTH option specifies the administrative authority level required to issue QUERY or SQL SELECT commands. By default any administrator can issue QUERY and SELECT commands. You can use this option to restrict the use of these commands.

- **RECLAIMDELAY**
This option delays the reclamation of a SnapLock volume, allowing remaining data to expire so that there is no need to reclaim the volume.
- **RECLAIMPERIOD**
This option allows you to set the number of days for the reclamation period of a SnapLock volume.
- **REORGBEGINTIME**
The REORGBEGINTIME option specifies the earliest time that the IBM Spectrum Protect server can start a table or index reorganization.
- **REORGDURATION**
The REORGDURATION option specifies an interval during which server-initiated table or index reorganization can start.
- **REPORTRETRIEVE**
The REPORTRETRIEVE option reports on restore or retrieve operations that are performed by client nodes or administrators. The default is NO.
- **REPLBATCHSIZE**
The REPLBATCHSIZE option specifies the number of client files that are to be replicated in a batch, within the same server transaction. This option affects only the node replication processes and works with the REPLSIZETHRESH option to improve node replication processing.
- **REPLSIZETHRESH**
The REPLSIZETHRESH option specifies, in megabytes, a threshold for the amount of data replicated, within the same server transaction.
- **REQSYSAUTHOUTFILE**
The REQSYSAUTHOUTFILE option specifies if system authority is required for administrative commands that cause IBM Spectrum Protect to write to an external file.
- **RESOURCETIMEOUT**
The RESOURCETIMEOUT option specifies how long the server waits for a resource before canceling the pending acquisition of a resource. When a timeout occurs the request for the resource will be canceled.
- **RESTHTTPSPORT**
The RESTHTTPSPORT option specifies the port number to be used for Hypertext Transfer Protocol Secure (HTTPS) communication between the Operations Center and the hub server.
- **RESTOREINTERVAL**
The RESTOREINTERVAL option specifies how long a restartable restore session can be saved in the server database. As long as the restore session is saved in the database, it can be restarted from the point at which it stopped.
- **RETENTIONEXTENSION**
The RETENTIONEXTENSION option specifies the number of days to extend the retention date of a SnapLock volume. This option allows the server to extend the retention date of a SnapLock volume in order to avoid excessive reclamation.
- **AIX Linux Windows SANDISCOVERY**
The SANDISCOVERY option specifies whether the IBM Spectrum Protect SAN discovery function is enabled.
- **AIX Linux Windows SANDISCOVERYTIMEOUT**
The SANDISCOVERYTIMEOUT option specifies the amount of time allowed for host bus adapters to respond when they are queried by the SAN discovery process. Once the time specified for the SANDISCOVERYTIMEOUT is reached, the process times out.
- **AIX Linux Windows SANREFRESHTIME**
The SANREFRESHTIME option specifies the amount of time that elapses before the cached SAN discovery information is refreshed. The SANREFRESHTIME option has a default value of 0, which means that there is no SAN discovery cache. The information is obtained directly from the host bus adapter (HBA) every time the server performs a SAN discovery operation.
- **SEARCHMPQUEUE**
The SEARCHMPQUEUE option specifies the order in which the server satisfies requests in the mount queue. If the option is specified, the server first tries to satisfy requests for volumes that are already mounted. These requests may be satisfied before other requests, even if the others have been waiting longer for the mount point. If this option is not specified, the server satisfies requests in the order in which they are received.
- **Windows SECUREPIPES**
When using the named pipes protocol, enabling SECUREPIPES forces the server to check the Windows group designated by ADSMGROUPNAME in order to authenticate a client node/user.
- **SERVERDEDUPTXNLIMIT**
The SERVERDEDUPTXNLIMIT option specifies the maximum size of objects that can be deduplicated on the server.
- **SHMPORT**
AIX Linux The SHMPORT option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection. **Windows** The SHMPORT option specifies the port that the server listens on for shared memory connections.
- **SHREDDING**
The SHREDDING option specifies whether shredding of deleted sensitive data is performed automatically or manually. Shredding applies only to data in storage pools that have been explicitly configured to support shredding.

- **SNMPHEARTBEATINTERVAL**
The SNMPHEARTBEATINTERVAL option specifies the interval in minutes between queries of the IBM Spectrum Protect server.
- **SNMPMESSAGECATEGORY**
The SNMPMESSAGECATEGORY option specifies the trap types used when messages are forwarded from the server, through the Simple Network Management Protocol (SNMP) subagent, to the SNMP manager.
- **SNMPSUBAGENT**
The SNMPSUBAGENT option specifies the parameters needed for the IBM Spectrum Protect subagent to communicate with the Simple Network Management Protocol (SNMP) daemon. This option is only to configure the SNMP subagent for communicating with the SNMP agent; it is ignored by the server.
- **SNMPSUBAGENTHOST**
The SNMPSUBAGENTHOST option specifies the location of the IBM Spectrum Protect Simple Network Management Protocol (SNMP) subagent. The default for this option is 127.0.0.1.
- **SNMPSUBAGENTPORT**
The SNMPSUBAGENTPORT option specifies the port number of the IBM Spectrum Protect Simple Network Management Protocol (SNMP) subagent.
- **SSLFIPSMODE**
The SSLFIPSMODE option specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for Secure Sockets Layer (SSL). The default is NO.
- **SSLINITTIMEOUT**
The SSLINITTIMEOUT option specifies the time, in minutes, that the server waits for a Secure Sockets Layer (SSL) session to complete initialization before the server cancels the session.
- **SSLTCPADMINPORT**
The SSLTCPADMINPORT option specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions only. The sessions are for the command-line administrative client.
- **SSLTCPPOINT**
The SSLTCPPOINT option specifies the Secure Sockets Layer (SSL) port number for SSL-enabled sessions only. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.
- **TCPADMINPORT**
The TCPADMINPORT option specifies the port number on which the server TCP/IP communication driver waits for requests for TCP/IP and SSL-enabled sessions other than client sessions. This includes administrative sessions, server-to-server sessions, storage agent sessions, library client sessions, managed server sessions, and event server sessions.
- | | |
|-----|-------|
| AIX | Linux |
|-----|-------|

TCPBUFSIZE
The TCPBUFSIZE option specifies the size of the buffer used for TCP/IP send requests. During a restore, client data moves from the IBM Spectrum Protect session component to a TCP communication driver. The TCPBUFSIZE option determines if the server sends the data directly from the session buffer or copies the data to the TCP buffer. A 32 KB buffer size forces the server to copy data to its communication buffer and flush the buffer when it fills.
- **TCPNODELAY**
The TCPNODELAY option specifies whether the server disables the delay of sending successive small packets on the network.
- **TCPPOINT**
The TCPPOINT option specifies the port number on which the server TCP/IP communication driver waits for requests for client sessions. The server TCP/IP communication driver listens on this port for both TCP/IP and SSL-enabled sessions from the client.
- **TCPWINDOWSIZE**
The TCPWINDOWSIZE option specifies, in kilobytes, the amount of receive data that can be buffered at one time on a TCP/IP connection. The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window lets the sender continue sending data, and may improve communication performance, especially on fast networks with high latency.
- **TECBEGINEVENTLOGGING**
The TECBEGINEVENTLOGGING option specifies whether event logging for the Tivoli® receiver should begin when the server starts up. If the TECHOST option is specified, TECBEGINEVENTLOGGING defaults to YES.
- **TECHOST**
The TECHOST option specifies the host name or IP address for the Tivoli event server.
- **TECPOINT**
The TECPOINT option specifies the TCP/IP port address on which the Tivoli event server is listening. This option is only required if the Tivoli event server is on a system that does not have a Port Mapper service running.
- **TECUTF8EVENT**
The TECUTF8EVENT option allows the IBM Spectrum Protect administrator to send information to the Tivoli Enterprise Console® (TEC) server in UTF-8 data format. The default is No. You can display whether or not this option is enabled by issuing the QUERY OPTION command.

- **THROUGHPUTDATATHRESHOLD**
The THROUGHPUTDATATHRESHOLD option specifies a throughput threshold that a client session must reach to prevent being canceled after the time threshold is reached.
- **THROUGHPUTTIMETHRESHOLD**
The THROUGHPUTTIMETHRESHOLD option specifies the time threshold for a session after which it may be canceled for low throughput.
- **Windows** **TIMEFORMAT**
The TIMEFORMAT option specifies the format in which time is displayed by the server.
- **TXNGROUPMAX**
The TXNGROUPMAX option specifies the number of objects that are transferred as a group between a client and the server between transaction commit points. The minimum value is 4 objects and the maximum value is 65000 objects. The default value is 4096 objects. The objects transferred are actual files, directories, or both. The server counts each file or directory as one object.
- **UNIQUETDPTECEVENTS**
The UNIQUETDPTECEVENTS option generates a unique Tivoli Enterprise Console (TEC) event class for each individual IBM Spectrum Protect message, including client, server, and IBM Spectrum Protect Data Protection client messages. The default is No.
- **UNIQUETECEVENTS**
The UNIQUETECEVENTS option generates a unique Tivoli Enterprise Console (TEC) event class for each individual IBM Spectrum Protect message. The default is No.
- **USEREXIT**
The USEREXIT option specifies a user-defined exit that will be given control to manage an event.
- **VERBCHECK**
The VERBCHECK option specifies that the server will do additional error checking on the structure of commands sent by the client. This option should only be enabled when the client sends incorrectly formed requests to the server, causing the server to crash. When this option is enabled, you will get a protocol error instead of a server crash.
- **VOLUMEHISTORY**
The VOLUMEHISTORY option specifies the name of files to be automatically updated whenever server sequential volume history information is changed. There is no default for this option.

Modifying server options

The server reads the server options file at server initialization. When you update a server option by editing the file, you must stop and start the server to activate the updated server options file.

About this task

You can change some options dynamically without stopping and starting the server, by using the SETOPT command. See SETOPT (Set a server option for dynamic update) for details.

AIX | **Linux** The dsmserv.opt.smp file (also provided at installation) contains the format of the options file and all the default settings. You can change any options in the dsmserv.opt.smp file. To have the server use the changed options, you must rename the file to dsmserv.opt. To activate an option within the server options file, remove the *>>> that precedes the option. The server ignores any options preceded by *>>>.

Windows You can modify server options by using the options file editor included in the IBM Spectrum Protect™ Console. This editor provides communications parameter detection, value validation, and help for all options. The options file editor is the preferred way to change server options, but you can also use a text editor.

Types of server options

Server options let you customize how some functions and processes work.

- **Server communication options**
You can use server options to specify server communication methods and their characteristics.
- **Server storage options**
IBM Spectrum Protect provides a number of options that you can specify to configure certain device and server storage operations.
- **Client-server options**
You can use server options to control client-server processing.

- Date, number, time, and language options
You can use server options to specify display formats for the dates, times, numbers, and national language.
- Database options
You can use server options to control some aspects of database processing.
- Data transfer options
You can use server options to control how IBM Spectrum Protect groups and transfers data.
- Message options
You can use server options to give you more flexibility in the way IBM Spectrum Protect issues messages.
- Event logging options
Options can help you manage event logging receivers.
- Security options and licensing options
You can use server options to customize server security and license audits.
- Miscellaneous options
You can use a variety of miscellaneous server options to customize IBM Spectrum Protect.

Server communication options

You can use server options to specify server communication methods and their characteristics.

Table 1. Communication options







Option	Description
ADMINCOMMTIMEOUT	The amount of time that the server waits for an administrative client message during an operation that causes a database update
ADMINIDLETIMEOUT	The amount of time an administrative client session can be idle
ADMINONCLIENTPORT	The port that determines whether administrative sessions can use the port specified in the TCPPORT option
COMMMETHOD	The server communication method
DBMTCPPOINT	The port number on which the TCP/IP communication driver for the database manager waits for client session requests
DNSLOOKUP	Control of use of Domain Name Services to lookup names of systems contacting the server
FIPSMODE	Specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for non-SSL operations.
LDAPCACHEDURATION	Determines the amount of time that authentication sessions, to the same node or administrator, are skipped. You might see a slight performance boost when skipping sessions.
LDAPURL	Specifies the LDAP directory server. Each setting must have the LDAP directory server name, a port number, and the base distinguished name of the namespace or suffix that the server maintains.
Windows NAMEDPIPENAME	Windows The named pipes communication method
NDMPCONTROLPORT	The internal communications port used for certain Network Data Management Protocol (NDMP) operations

Option	Description
NDMPENABLEKEEPALIVE	The TCP keepalive mechanism
AIX Linux Windows NDMPKEEPIDLEMINUTES	AIX Linux Windows The amount of idle time before the first TCP keepalive packet is sent
Windows NPBUFFERSIZE	Windows The size of the Named Pipes communication buffer
SHMPORT	AIX Linux The TCP/IP port address of a server when using shared memory Windows The port that the server listens on for shared memory connections
SNMPHEARTBEATINTERVAL	The interval in minutes between queries of the IBM Spectrum Protect server
SNMPMESSAGECATEGORY	The trap types used when messages are forwarded from the server
SNMPSUBAGENT	The parameters needed for the IBM Spectrum Protect subagent to communicate with the SNMP daemon
SNMPSUBAGENTHOST	The location of the IBM Spectrum Protect SNMP subagent
SNMPSUBAGENTPORT	The port address of the IBM Spectrum Protect SNMP subagent
SSLFIPSMODE	Specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for Secure Sockets Layer (SSL)
SSLTCPADMINPORT	The port address on which the server's TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client
SSLTCPPOINT	The SSL-only port number on which the server's TCP/IP communication driver waits for requests for SSL-enabled sessions from the following sources: <ul style="list-style-type: none"> • Command line backup-archive client • Backup-archive GUI • Administrative client • Application programming interface (API)
TCPADMINPORT	The TCP/IP port number for administrative sessions
AIX Linux TCPBUFSIZE	AIX Linux The size of the buffer used for TCP/IP send requests
TCPPOINT	The TCP/IP port number for client sessions
TCPWINDOWSIZE	The client node TCP/IP sliding window

Server storage options

IBM Spectrum Protect™ provides a number of options that you can specify to configure certain device and server storage operations.

Table 1. Server storage options

Option	Description
3494SHARED	Enables sharing of a 3494 library with applications other than IBM Spectrum Protect.
ACSACCESSID	The ID for the ACS access control.
ACSLCKDRIVE	Allows the drives within the ACSLS libraries to be locked.
ACSQUICKINIT	Allows a quick or full initialization of the ACSLS library.
ACSTIMEOUTX	The multiple for the built-in timeout value for the ACSLS API.
ASSISTVCRRECOVERY	Specifies whether the server assists an IBM 3590 drive in recovering from a lost or corrupted Vital Cartridge Records (VCR) condition.
CHECKTAPEPOS	Specifies whether the server validates data position on tape.
CLIENTDEDUPTXNLIMIT	Specifies the maximum size of a transaction when client-side deduplicated data is backed up or archived.
DEDUPREQUIRESBACKUP	Specifies whether volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and whether duplicate data can be discarded before the storage pools are backed up.
DEDUPTIER2FILESIZE	File size at which Tier 2 processing is used for data deduplication.
DEDUPTIER3FILESIZE	File size at which Tier 3 processing is used for data deduplication.
DEVCONFIG	The name of the file that store backup copies of device configuration information.
DRIVEACQUIRERETRY	The number of times that the server retries the acquisition of a drive in an IBM 349x library that is shared among multiple applications.
ENABLENASDEDUP	Specifies whether the server deduplicates data that is stored by a NetApp network-attached storage (NAS) file server.
NUMOPENVOLSALLOWED	The number of input FILE volumes in a deduplicated storage pool that can be open at one time.
RECLAIMDELAY	The number of days that the reclamation of a SnapLock volume is delayed.
RECLAIMPERIOD	The number of days for the reclamation period of a SnapLock volume
RESOURCETIMEOUT	The length of time that the server waits for a resource before canceling the pending acquisition of the resource.
RETENTIONEXTENSION	The number of days to extend the retention date of a SnapLock volume.
 SANDISCOVERY	 Whether the IBM Spectrum Protect SAN discovery function is enabled.
 SANDISCOVERYTIMEOUT	 Amount of time before the SAN discovery process times out.
 SANREFRESHTIME	 Amount of time before cached SAN discovery information is refreshed.
SEARCHMPQUEUE	The order in which the server satisfies requests in the mount queue.
SERVERDEDUPTXNLIMIT	Specifies the maximum size of objects that can be deduplicated on the server.

Client-server options

You can use server options to control client-server processing.

Table 1. Client-Server options

Option	Description
--------	-------------

Option	Description
COMMTIMEOUT	The number of seconds the server waits for a response from a client before timing out the client session
DISABLESCHEDS	Whether administrative and client schedules are disabled during the IBM Spectrum Protect server recovery scenario
IDLETIMEOUT	The number of minutes the server allows a client session to remain idle before timing out the client session
MAXSESSIONS	The maximum number of simultaneous client sessions with the server
THROUGHPUTDATATHRESHOLD	The throughput threshold that a client session must reach to prevent being canceled after the time threshold is reached
THROUGHPUTTIMETHRESHOLD	The time threshold for a session after which it may be canceled for low throughput
VERBCHECK	Whether additional error checking is done for commands sent by the client

Date, number, time, and language options

You can use server options to specify display formats for the dates, times, numbers, and national language.

Table 1. Date, number, time, and language options

Option	Description
Windows DATEFORMAT	Windows The format by which dates are displayed
LANGUAGE	The national language is used to present client messages
Windows NUMBERFORMAT	Windows The format for displaying numbers
Windows TIMEFORMAT	Windows The format displaying times

Database options

You can use server options to control some aspects of database processing.

Table 1. Database options

Option	Description
ACTIVELOGDIRECTORY	The new directory for the location where the active log is stored. Use this option to change the location of the active log.
ACTIVELOGSIZE	The maximum size of the active log.
ALLOWREORGINDEX	Server-initiated index reorganization.
ALLOWREORGTABLE	Server-initiated table reorganization.
ARCHLOGDIRECTORY	The directory that the database manager can archive a log file into after all the transactions represented in that log file are completed.
ARCHFAILOVERLOGDIRECTORY	The directory in which the server tries to store archive log files that cannot be stored in the archive log directory.
DBDIAGLOGSIZE	The maximum size of the database manager diagnostic log files.
DBDIAGPATHFSTHRESHOLD	The threshold for free space on the file system or disk that contains the database manager diagnostic log files.
DBMEMPERCENT	The percentage of system memory that is dedicated to the database.
DISABLEREORGTABLE	Disables table reorganization for specific tables.
FSUSEDTHRESHOLD	The percentage of the file system that can be used by the database before an alert message is issued.
MIRRORLOGDIRECTORY	The directory for mirroring the active log path.

Option	Description
REORGBEGINTIME	The earliest time that the IBM Spectrum Protect server can start a table or index reorganization.
REORGDURATION	The interval during which server-initiated table or index reorganization can start.

Data transfer options

You can use server options to control how IBM Spectrum Protect™ groups and transfers data.

Table 1. Group options

Option	Description
MOVEBATCHSIZE	The number of files that are to be moved and grouped in a batch, within a transaction
MOVESIZETHRESH	The threshold for the amount of data moved as a batch, within the same server transaction
NDMPPORTRANGE	The IP address associated with the interface in which the server receives all Network Data Management Protocol (NDMP) backup data
NDMPREFDATAINTERFACE	The IP address associated with the interface in which the server receives all Network Data Management Protocol (NDMP) backup data
REPLBATCHSIZE	The number of files that are to be replicated in a batch, within the same server transaction
REPLSIZETHRESH	The threshold for the amount of data replicated as a batch, within the same server transaction
TXNGROUPMAX	The number of files that are transferred as a group between a client and the server between transaction commit points

Message options

You can use server options to give you more flexibility in the way IBM Spectrum Protect™ issues messages.

Table 1. Message options

Option	Description
EXPQUIET	Whether IBM Spectrum Protect sends detailed informational messages during expiration processing
MESSAGEFORMAT	Whether a message number is displayed in all lines of a multi-line message
MSGINTERVAL	The time, in minutes, between messages prompting an operator to mount a tape for IBM Spectrum Protect

Event logging options

Options can help you manage event logging receivers.

Table 1. Event logging options

Option	Description
EVENTSERVER	Whether the server should try to contact the event server when the server starts up
FILEEXIT	A file to which enabled events are routed (binary format)
FILETEXTEXIT	A file to which enabled events are routed (readable format)
REPORTRETRIEVE	Record client restore and retrieve operations
TECBEGINEVENTLOGGING	Whether event logging for the TIVOLI receiver should begin when the server starts up

Option	Description
TECHOST	The host name or IP address for the Tivoli Enterprise Console (TEC) event server
TECPORT	The TCP/IP port address on which the Tivoli Enterprise Console event server is listening
TECUTF8EVENT	A Tivoli Enterprise Console event sent from the IBM Spectrum Protect server in UTF8 format
UNIQUETDPTECEVENTS	Events from an IBM Spectrum Protect Data Protection client that are sent to the Tivoli Enterprise Console as unique events
UNIQUETECEVENTS	Events sent to the Tivoli Enterprise Console as unique
USEREXIT	A user-defined exit that will be given control to manage an event

Security options and licensing options

You can use server options to customize server security and license audits.

Table 1. Security and licensing options

Option	Description
Windows ADMSGROUPNAME	Windows The name of a Windows group
AUDITSTORAGE	Specifies that during a license audit operation, the server calculates, by node, the amount of backup, archive, and space management storage in use
BACKUPINITIATIONROOT	Specifies whether the server overrides node parameter values for users who are not IBM Spectrum Protect authorized users
LDAPURL	Specifies the LDAP directory server. Each setting must have the LDAP directory server name, a port number, and the base distinguished name of the namespace or suffix that the server maintains.
Windows NPAUDITFAILURE	Windows Specifies that a node can access only its own data
Windows NPAUDITSUCCESS	Windows Specifies that an event is sent to the event log when a client node user is authenticated for access to the server through SECUREPIPE
QUERYAUTH	The administrative authority level required to issue QUERY or SQL SELECT commands
REQSYSAUTHOUTFILE	Specifies if system authority is required for administrative commands that cause IBM Spectrum Protect to write to an external file
Windows SECUREPIPES	Windows With named pipes protocol, specifies that the server checks the Windows group to authenticate a client
SHREDDING	Specifies whether shredding of deleted sensitive data is done automatically or manually

Related reference:

Server communication options

Miscellaneous options

You can use a variety of miscellaneous server options to customize IBM Spectrum Protect™.

Table 1. Miscellaneous options

Option	Description
ALIASHALT	Allows administrators to give the IBM Spectrum Protect HALT command a different name

Option	Description
DISPLAYLFINFO	Specifies whether accounting records and summary table entries report the storage agent name
EXPINTERVAL	The interval between automatic inventory expiration processes
FFDCLOGNAME	The name for the first failure data capture (FFDC) log
FFDCMAXLOGSIZE	The maximum size of the first failure data capture (FFDC) log
NOPREEMPT	Specifies that no operation can preempt another for access to a volume and that only a database backup operation can preempt another operation for access to a device
NORETRIEVEDATE	Specifies that the server does not update the retrieve date of a file in a disk storage pool when a client restores or retrieves the file
RESTOREINTERVAL	The length of time that a restartable restore session can be saved in the server database
VOLUMEHISTORY	The name of the file to be automatically updated whenever server sequential volume history information is changed

3494SHARED

The 3494SHARED option specifies whether an IBM® 3494 library can share applications other than IBM Spectrum Protect™.

The default is NO, meaning that no application other than IBM Spectrum Protect can share the 3494. When you set this option to YES, for every mount request, IBM Spectrum Protect determines if each drive is in use. After the query completes, IBM Spectrum Protect selects an available drive that is not in use by another application. Enable sharing only if you have more than two drives in your library. If you are currently sharing an IBM 3494 library with other applications, you must specify this option.

Syntax

```
>>-3494SHARED--+-Yes-+----->>
          '-No--'
```

Parameters

- Yes
Specifies that other applications can share the 3494 library.
- No
Specifies that no other applications can share the 3494 library.

Examples

Enable sharing of a 3494 library:

```
3494shared yes
```

ACSACCESSID

The ACSACCESSID option specifies the ID for the ACS access control for an ACSLS library.

Syntax

```
>>-ACSACCESSID--name-----<<
```

Parameters

name

Specifies a 1 to 64 character ID. The default ID is your local host name.

Examples

```
acsaccessid region
```

ACSLOCKDRIVE

The ACSLOCKDRIVE option specifies if the drives within the ACSLS libraries are locked. Drive locking ensures the exclusive use of the drive in the ACSLS library in a shared environment. However, there is some performance gain if libraries are not locked. When other applications do not share the IBM Spectrum Protect™ drives, drive locking is not required.

Syntax

```
>>-ACSLOCKDRIVE---+-Yes-+-----<<  
                    '-No--'
```

Parameters

Yes

Specifies that drives are locked.

No

Specifies that drives are not locked.

Examples

```
acslockdrive yes
```

ACSQUICKINIT

The ACSQUICKINIT option specifies whether, at server startup, the initialization of the ACSLS library is a quick or full initialization. The default is Yes. A quick initialization avoids the overhead associated with synchronizing the IBM Spectrum Protect™ server inventory with the ACSLS library inventory (through an audit of the library).

Syntax

```
>>-ACSQUICKINIT---+-Yes-+-----<<  
                    '-No--'
```

Parameters

Yes

Specifies that a quick initialization of the ACSLS library is performed. When the option is set to Yes, IBM Spectrum Protect bypasses library inventory verification, initializing the library quickly, and making it available to IBM Spectrum Protect sooner than if a full initialization is done.

This option should be set to Yes when it is known that the physical library inventory and the IBM Spectrum Protect library inventory have not changed and an audit is not needed.

No

Specifies that a full initialization of the ACSLS library and library inventory is performed. When the option is set to No, IBM Spectrum Protect synchronizes its library volume inventory with what is reported by the ACSLS library manager.

Examples

```
acsquickinit yes
```

ACSTIMEOUTX

The ACSTIMEOUTX option specifies the multiple for the built-in timeout value for ACSLS APIs. The built-in timeout value for the ENTER, EJECT, and AUDIT ACS API is 1800 seconds; for all other ACSLS APIs it is 600 seconds. For example, if the multiple value specified is 5, the timeout value for audit API becomes 9000 seconds, and all other APIs become 3000 seconds.

Syntax

```
>>-ACSTIMEOUTX--value-----<<
```

Parameters

value

Specifies the multiple for the built-in timeout value for ACSLS API. The range is from 1 to 100. The default is 1.

Examples

```
acstimeoutx 1
```

ACTIVELOGDIRECTORY

The ACTIVELOGDIRECTORY option specifies the name of the directory where all active logs are stored.

This option is appended to the options file when the DSMSERV FORMAT command is run. Under normal operating conditions, the option does not need to be changed. See DSMSERV FORMAT (Format the database and log) for guidance on this option.

Syntax

```
>>-ACTIVELOGDirectory--dir_name-----<<
```

Parameters

dir_name

Specifies a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. If you change the active log directory, IBM Spectrum Protect™ moves the existing active logs to the location that is specified by this directory. The maximum number of characters is 175.

Examples

AIX | Linux

```
activelogdirectory /tsm/activelogdir
```

Windows

```
activelogdirectory c:\tsmserv1\activelogdir
```

ACTIVELOGSIZE

The ACTIVELOGSIZE option sets the total log size.

This option is appended to the options file when the DSMSERV FORMAT command is run. Under normal operating conditions the option does not need to be changed. See DSMSERV FORMAT (Format the database and log) for guidance on this option.

Syntax

```
.-16GB-----.  
>>-ACTIVELOGSize---+megabytes-+-----<<
```

Parameters

megabytes

Specifies the size of the active log file in megabytes. The minimum value is 2048 MB (2 GB); the maximum is 524,288 MB (512 GB). If an odd number is specified, the value is rounded up to the next even number. The default is 16,384 MB (16 GB).

The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

Examples

```
activelogsiz 8192
```

ADMINCOMMTIMEOUT

The ADMINCOMMTIMEOUT option specifies how long the server waits for an expected administrative client message during an operation that causes a database update.

If the length of time exceeds this time-out period, the server ends the session with the administrative client. You may want to increase the time-out value to prevent administrative client sessions from timing out.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
.-60-----.  
>>-ADMINCOMMTimeout---+seconds-+-----<<
```

Parameters

seconds

Specifies the maximum number of seconds that a server waits for an administrative client response. The default value is 60. The minimum value is 1.

Examples

```
admincommtimeout 60
```


ADMINIDLETIMEOUT

The ADMINIDLETIMEOUT option specifies the amount of time, in minutes, that an administrative client session can be idle before the server cancels the session.

If there is a heavy network load in your environment, you might want to increase the time-out value to prevent administrative clients from timing out. However, a large number of idle sessions could prevent other users from connecting to the server.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
                .-15-----.  
>>-ADMINIDLETIMEOUT--+-minutes+-----><
```

Parameters

minutes

Specifies the maximum number of minutes that a server waits for an idle administrative client. The default value is 15 minutes. The minimum value is 1 minute.

Examples

```
adminidletimeout 20
```

ADMINONCLIENTPORT

The ADMINONCLIENTPORT option specifies whether the TCPPOINT can be used by administrative sessions. The default is YES.

Syntax

```
>>-ADMINONCLIENTPORT--+-YES-+-----><  
                '-NO--'
```

Parameters

YES

If the option is set to YES, or if the TCPPOINT and TCPADMINPORT are the same value (the default), administrative sessions can use the TCPPOINT.

NO

If the option is set to NO, and if the TCPADMINPORT value is different than the TCPPOINT value, administrative sessions cannot use the TCPPOINT.

Examples

Specify that the TCPPOINT can be used by administrative sessions.

```
adminonclientport yes
```

Windows

ADSMGROUPNAME

The ADSMGROUPNAME option specifies the name of a Windows group. A client node must be a member of this group to use the IBM Spectrum Protect™ server through NT Unified Logon. The client node must also be a registered IBM Spectrum Protect client node.

Syntax

```
>>-ADSMGROUPname--group_name-----<<
```

Parameters

group_name
Specifies a Windows group name.

Examples

Specify IDD as a Windows group:

```
adsmgroup idd
```

ALIASHALT

The ALIASHALT option allows administrators to give the IBM Spectrum Protect™ **HALT** command a different name.

The administrative client recognizes an alias for the HALT command when the client is started with the CHECKALIASHALT option specified. See Administrative client options for details.

Syntax

```
>>-ALIASHALT--newname-----<<
```

Parameters

newname
Specifies the alias of the HALT command for shutting down the IBM Spectrum Protect server. Minimum length of *newname* is 1; maximum length is 16.

Examples

```
aliashalt tsmhalt
```

ALLOWDESAUTH

The ALLOWDESAUTH option specifies whether to allow use of the Data Encryption Standard (DES) algorithm for authentication between a server and a backup-archive client.

To prevent the use of DES, specify a value of NO for the ALLOWDESAUTH option.

To configure the IBM Spectrum Protect™ server to be in compliance with the NIST SP800-131A standard, set this option to NO. Restrictions:

- The backup-archive client must be running Version 6.3 or later if you authenticate to a server with the ALLOWDESAUTH option set to NO.
- Automatic deployment of the backup-archive client fails if this option is set to NO.

Syntax

```
.-ALLOWDESAUTH--Yes-----  
>>-+-----<<  
'-ALLOWDESAUTH---No---'  
  '-Yes-'
```

Parameters

- Yes
Specifies that the server allows authentication with any backup-archive clients that use DES-based encryption. The default is YES.
- No
Specifies that the server rejects any backup-archive clients that attempt to authenticate with DES-based encryption.

Examples

Specify that the server rejects any backup-archive clients that attempt to authenticate with DES encryption:

```
allowdesauth no
```

Specify that the server allows authentication with any backup-archive clients that use DES encryption:

```
allowdesauth yes
```

ALLOWREORGINDEX

The ALLOWREORGINDEX option specifies whether server-initiated index reorganization is enabled or disabled.

The default is YES.

Syntax

```
>>-ALLOWREORGINDEX---+-Yes-+-----<-----<<  
                        '-No--'
```

Parameters

- Yes
Specifies that server-initiated index reorganization is enabled.
- No
Specifies that server-initiated index reorganization is disabled.

Example

Specify that server-initiated index reorganization is enabled.

```
allowreorgindex yes
```

ALLOWREORGTABLE

The ALLOWREORGTABLE option specifies whether server-initiated table reorganization is enabled or disabled.

The default is YES.

Syntax

```
>>-ALLOWREORGTABLE---+-Yes-+-----<-----<<  
                        '-No--'
```

Parameters

- Yes
Specifies that server-initiated table reorganization is enabled.
- No
Specifies that server-initiated table reorganization is disabled.

Examples

Specify that server-initiated table reorganization is disabled.

```
allowreorgtable no
```

ARCHFAILOVERLOGDIRECTORY

The ARCHFAILOVERLOGDIRECTORY option specifies the directory which the server uses to store archive log files that cannot be stored in the archive log directory.

This option is appended to the options file when the DSMSERV FORMAT command is run. Typically the directory does not need to be changed.

Syntax

```
>>-ARCHFailoverlogdirectory--dir_name-----<<
```

Parameters

dir_name

Specifies a fully qualified directory name. The maximum number of characters is 175.

Examples

AIX | Linux

```
archfailoverlogdirectory /tsm/archfailoverlog
```

Windows

```
archfailoverlogdirectory c:\tsmserv1\archfailoverlog
```

ARCHLOGCOMPRESS

You can enable or disable compression of archive logs on the IBM Spectrum Protect™ server. By compressing the archive logs, you reduce the amount of space that is required for storage.

The ARCHLOGCOMPRESS server option specifies whether log files that are written to the archive directory for logs are compressed.

Syntax

```
>>-ARCHLOGCOMPRESS--.-No--.
                    +-----+-----<<
                    '-Yes-'
```

Parameters

No

Specifies that log files that are written to the archive log directory are not compressed. The default is No.

Yes

Specifies that log files that are written to the archive log directory are compressed.

Restriction: Use caution when you enable the ARCHLOGCOMPRESS server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log directory file system usage nears out of space conditions, the ARCHLOGCOMPRESS server option must be disabled. You can use the SETOPT command to disable archive log compression immediately without halting the server.

Example

To enable compression of log files that are written to the archive log directory, specify the following option:

```
archlogcompress yes
```

ARCHLOGDIRECTORY

The ARCHLOGDIRECTORY option specifies a directory that the database manager can archive a log file into after all the transactions represented in that log file are completed.

This option is appended to the options file when the DSMSERV FORMAT command is run.

Syntax

```
>>-ARCHLOGDirectory--dir_name-----<<
```

Parameters

dir_name

Specifies a fully qualified directory name. The maximum number of characters is 175.

Examples

AIX | Linux

```
archlogdirectory /tsm/archlog
```

Windows

```
archlogdirectory d:\tmserv1\archlog
```

ARCHLOGUSEDTHRESHOLD

The ARCHLOGUSEDTHRESHOLD option specifies when to start an automatic database backup in relation to the percentage of archive log file space used. The default is 80 percent.

The ARCHLOGUSEDTHRESHOLD option prevents frequent automatic backups. For example, if the archive log file directory resides on a file system or drive that is 400 GB, a database backup is triggered if there is less than 80 GB of free space. Repeated database backups might cause the server to use an excessive amount of scratch tapes.

Syntax

```
>>-ARCHLOGUSEDTHRESHOLD--+-value-+-----<<
```

Parameters

value

The percentage of archive log file space used before an automatic backup starts.

Specify to start an automatic backup when 90 percent of archive log file space is used.

```
archlogusedthreshold 90
```

ASSISTVCRRECOVERY

The ASSISTVCRRECOVERY option specifies whether IBM Spectrum Protect™ assists an IBM® 3590 drive in recovering from a lost or corrupted Vital Cartridge Records (VCR) condition. If you specify YES (the default) and if IBM Spectrum Protect detects an error

during the mount processing, it locates to the end-of-data during the dismount processing to allow the drives to restore the VCR. During the tape operation, there might be some small effect on performance because the drive cannot complete a fast locate with a lost or corrupted VCR. However, there is no loss of data.

Syntax

```
>>-ASSISTVCRREcovery--+-Yes-+-----><
                    '-No--'
```

Parameters

Yes
Specifies server assistance in recovery.

No
Specifies no server assistance in recovery.

Examples

Turn off recovery assistance:

```
assistvcrrecovery no
```

AUDITSTORAGE

As part of a license audit operation, the server calculates, by node, the amount of server storage used for backup, archive, and space-managed files. For servers managing large amounts of data, this calculation can take a great deal of CPU time and can stall other server activity. You can use the AUDITSTORAGE option to specify that storage is not to be calculated as part of a license audit.

Note: This option was previously called NOAUDITSTORAGE.

Syntax

```
>>-AUDITStorage---+-Yes-+-----><
                    '-No--'
```

Parameters

Yes
Specifies that storage is to be calculated as part of a license audit. The default is Yes.

No
Specifies that storage is not to be calculated as part of a license audit.

Examples

```
auditstorage yes
```

BACKUPINITIATIONROOT

The BACKUPINITIATIONROOT option specifies whether the server overrides node parameter values for users who are not IBM Spectrum Protect™ authorized users.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-BACKUPINITIATIONROOT--+--ON--+----->>
      '-OFF-'
```

Parameters

ON

Specifies that sessions from clients on AIX®, Linux, Mac OS X, and Solaris operating systems, where the users are not IBM Spectrum Protect authorized users, are prevented from initiating backup operations. This is the default. The server overrides the value for the BACKUPINITIATION parameter that is specified in the REGISTER NODE and UPDATE NODE commands.

Tip: For an overview of IBM Spectrum Protect authorized users, see UNIX and Linux client root and authorized user tasks.

OFF

Specifies that the node value for the BACKUPINITIATION parameter is used. The BACKUPINITIATION parameter is specified in the REGISTER NODE and UPDATE NODE commands.

Example

Specify that the node value for the BACKUPINITIATION parameter is used.

```
backupinitiationroot off
```

CHECKTAPEPOS

The CHECKTAPEPOS option specifies whether the IBM Spectrum Protect™ server validates the position of data blocks on tape.

The CHECKTAPEPOS option applies only to operations that use tape drives. It does not apply to non-tape, sequential-access device classes such as FILE. If the server information about position does not match the position that is detected by the drive, an error message is displayed, the transaction is rolled back, and the data is not committed to the database.

Using the CHECKTAPEPOS option, you can enable append-only mode for IBM® LTO Generation 5 and later drives, and for any drives that support this feature. When it is enabled, the drive issues an error after it receives instructions to overwrite any data on the currently mounted volume. The IBM Spectrum Protect server repositions the tape to the correct block and continues writing data. Append-only mode provides added protection by preventing most data overwrite situations. If you are using a drive that supports this feature, you can validate data position on tape by using both IBM Spectrum Protect and the drive or you can enable one or the other.

Note: When you use SAN Tape acceleration functions in the fabric or SAN switch, set the CHECKTAPEPOS option to DRIVEonly or No to avoid false positive positioning errors. The IBM Spectrum Protect CHECKTAPEPOS server option does not require an append-only capable drive.

Changes to the CHECKTAPEPOS option affect mounts only after the update to the drive is complete.

The default is YES.

Syntax

```
>>-CHECKTAPEPOS--+--Yes----->>
      +-No-----+
      +-TSMonly----+
      '-DRIVEonly-'
```

Parameters

Yes

Specifies that the IBM Spectrum Protect server validates data position on tape. For drives that support append-only mode, this parameter specifies that IBM Spectrum Protect enables the drive to also validate the data position during each WRITE operation to prevent data overwrite. Yes is the default.

No

Specifies that all data position validation is turned off.

TSMonly

Specifies that the IBM Spectrum Protect server validates data position on tape. The server does not use append-only mode even if the drive supports the feature.

DRIVEonly

Specifies that the IBM Spectrum Protect server enables append-only mode for drives that support this feature. The server does not validate the data position on tape.

Example

Validate data position on tape and enable append-only mode for a supported drive:

```
checktapepos yes
```

CLIENTDEDUPTXNLIMIT

The CLIENTDEDUPTXNLIMIT option specifies the maximum size of a transaction when client-side deduplicated data is backed up or archived.

When you use client-side deduplication for large objects, intensive database activity can result from long-running transactions that are required to update the database. High levels of database activity can produce the following symptoms:

- Reduced throughput for client backup and archive operations
- Resource contention resulting from concurrent server operations
- Excessive recovery log activity

The extent to which these symptoms occur depends on the number and size of objects being stored using client-side data deduplication, the intensity and type of concurrent operations taking place on the IBM Spectrum Protect™ server, and the IBM Spectrum Protect server configuration.

With the CLIENTDEDUPTXNLIMIT server option, you can specify a maximum size, in gigabytes, for transactions when client-side deduplicated data is backed up or archived. If an object or set of objects in a single transaction exceeds the limit specified by CLIENTDEDUPTXNLIMIT, the objects are not deduplicated by the client, and the transaction can fail. You can specify a value 32 - 102400 GB. The default value is 5120 GB.

If an object or set of objects in a single transaction exceeds the limit specified by CLIENTDEDUPTXNLIMIT, the objects or set of objects is not deduplicated by the client. However, the objects are sent to the server. These objects can be deduplicated on the server, depending on whether the destination storage pool is configured for data deduplication and on the value of the SERVERDEDUPTXNLIMIT option. Objects in a deduplication-enabled storage pool that are less than the value of the SERVERDEDUPTXNLIMIT are deduplicated by a server duplicate-identification process.

The appropriate value for this option depends on the IBM Spectrum Protect server configuration and concurrent server activity. You can specify a high value for this option if you minimize resource contention. To minimize resource contention, perform operations, such as backup, archive, duplicate identification (the IDENTIFY DUPLICATES command), and reclamation, at different times.

To update this server option without stopping and restarting the server, use the SETOPT command.

Syntax

```
                .-5120-----.  
>>-CLIENTDEDUPTXNlimit--+-gigabytes-+-----><
```

Parameters

gigabytes

Specifies the maximum size, in gigabytes, of objects that can be backed up or archived using client-side data deduplication. You can specify a value 32 - 102400. The default value is 5120.

Examples

Disable client-side data deduplication for all objects over 80 GB:

```
clientdeduptxnlimit 80
```


CLIENTDEPLOYCATALOGURL

The CLIENTDEPLOYCATALOGURL option specifies the location of the catalog file that is used for automatic client deployment operations.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-CLIENTDEPLOYCATalogurl----->
      .-https://public.dhe.ibm.com/storage/tivoli-storage-management/catalog/client/catalog.json-.
>--+url-----+><
```

Parameters

url

Specifies the URL from which the server downloads the catalog file for automatic client deployment operations. The catalog file stores properties for client deployment operations, including the location of the deployment packages. The default URL is `https://public.dhe.ibm.com/storage/tivoli-storage-management/catalog/client/catalog.json`.

To specify that the catalog file is downloaded from another location, use the SETOPT command to specify a custom URL. To reset the URL to the default value, issue the SETOPT command with an empty string: `""`. If you specify a custom URL, the custom URL is retained after the server is upgraded.

Example

Specify a custom URL of `https://customAddress`.

```
setopt clientdeploycatalogurl https://customAddress
```

Example

Restore the value of the CLIENTDEPLOYCATALOGURL option to the default.

```
setopt clientdeploycatalogurl ""
```

CLIENTDEPLOYUSELOCALCATALOG

The CLIENTDEPLOYCATALOGURL option specifies whether the local version of the catalog file is used for automatic client deployment operations.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-CLIENTDEPLOYUSELOCALcatalog--+-Yes+-----><
      .-No--.
```

Parameters

No

Specifies that the local version of the catalog file is not used. Instead, the catalog file is downloaded from the location that is specified by the CLIENTDEPLOYCATALOGURL option. The default value is NO.

Yes

Specifies that the local version of the catalog file is used. Catalog files are not downloaded during client deployment operations. If you set this option to YES, the value is retained after the server is upgraded.

Example

Specify that the local version of the catalog file is used.

```
setopt clientdeployuselocalcatalog yes
```

COMMMETHOD

The COMMMETHOD option specifies a communication method to be used by the server.

You can configure the server to use multiple communication methods. The more commonly used are the TCPIP, V6TCPIP, and SHAREDMEM communication methods. To specify multiple communication methods, enable each method by adding a COMMMETHOD stanza to the dsmserv.opt options file.

Important: When you enable a communication method, you must also add the options that are specific to the communication method to the options file.

Syntax

```
      .-TCPIP-----.  
>>-COMMMethod--+-NAMEDPIPE+-----><  
      +-NONE-----+  
      +-SHAREDMEM--+  
      +-SNMP-----+  
      +-TCPIP-----+  
      '-V6TCPIP---'
```

Parameters

You can choose one of the following communication methods:

Windows NAMEDPIPES

Windows Specifies the named pipes communication method option.

NONE

Specifies that no communication method is used. This option does not allow users to connect to the server and is useful for experimenting with policy commands.

SHAREDMEM

Specifies the shared memory communication method option. This method uses the same area of memory to send data between several applications at the same time. Both the server and the backup-archive client must be configured to support the shared memory communication method, and they must be installed on the same computer.

SNMP

Specifies the SNMP communication method option.

TCPIP

Specifies the TCP/IP communication method option. This option is the default. When TCPIP is specified, TCP/IP Version 4 is used exclusively.

V6TCPIP

Specifies the TCP/IP communication method option. If TCP/IP Version 4 and Version 6 are both configured, IBM Spectrum Protect™ uses both protocols simultaneously. If both COMMMETHOD TCPIP and COMMMETHOD V6TCPIP are specified, V6TCPIP overrides the specification of TCPIP. A valid domain name server (DNS) environment must be present to use either TCP/IP V4 or TCP/IP V6 if this option is specified.

Examples

Example of specifying multiple communication methods to be used by the server (TCP/IP and TCP/IP Version 6):

```
commmethod tcpip  
commmethod v6tcpip
```

COMMTIMEOUT

The COMMTIMEOUT option specifies how long the server waits for an expected client message during an operation that causes a database update. If the length of time exceeds this time-out, the server ends the session with the client. You may want to increase the time-out value to prevent clients from timing out. Clients may time out if there is a heavy network load in your environment or they are backing up large files.

The COMMTIMEOUT server option is used for non-administrative sessions. See the ADMINCOMMTIMEOUT option for administrative client sessions.

You can update this server option without stopping and restarting the server by using the SETOPT command.

Syntax

```
                .-60-----.  
>>-COMMTIMEOUT--+-seconds+-----><
```

Parameters

seconds

Specifies the maximum number of seconds that a server waits for a client response. The default value is 60. The minimum value is 1.

Examples

```
commtimeout 60
```

AIX

Linux

Windows

CONTAINERRESOURCETIMEOUT

The CONTAINERRESOURCETIMEOUT option specifies how long the server waits to complete a data store operation to a container storage pool.

Syntax

When a timeout occurs, any data that was stored in the container storage pool remains there. The data store operation ends, and the request for the container resource is canceled.

```
                .-180-----.  
>>-CONTAINERRESOURCETIMEOUT--+-minutes+-----><
```

Parameters

minutes

Specifies the maximum number of minutes that a server waits before an operation is canceled. The default value is 180 minutes. The minimum value is 1 minute.

Example

Specify that the server waits for 4 hours before a data store operation to a container storage pool is canceled.

```
containerresourcetimeout 240
```

Windows

DATEFORMAT

The DATEFORMAT option specifies the format in which dates are displayed by the server.

The DATEFORMAT value is overridden by the locale format if the locale is initialized at server startup. The locale is specified in the LANGUAGE option.

Syntax

```
>>-DATEformat--n-----><
```

Parameters

n

Select a number from 1 to 5 to identify the date format used by the server. The default value is 1.

1	MM/DD/YYYY
2	DD-MM-YYYY
3	YYYY-MM-DD
4	DD.MM.YYYY
5	YYYY.MM.DD

Examples

```
dateformat 4
```

DBDIAGLOGSIZE

This option helps to control the amount of space that is used by diagnostic log files.

The database manager uses diagnostic log files to log messages. You must control the size of the log files so that they do not fill the file system. Use the DBDIAGLOGSIZE option to set the amount of space that is used by the log files.

If you set a value in the range 2 - 9999, a maximum of 10 rotating diagnostic log files are retained. Each file name indicates the order in which the file was created. After a file is full, the next file is created. When the 10th file is full, the oldest file is deleted, and a new file is created. The following example shows how the rotating log files might look:

```
db2diag.14.log, db2diag.15.log, ... , db2diag.22.log, db2diag.23.log
```

When db2diag.23.log is full, db2diag.14.log is deleted, and db2diag.24.log is created.

The server checks the file space that contains the diagnostic log files every hour. Messages are displayed every 12 hours if either of the following conditions occur:

- The available space in the file system where the diagnostic log files are located is less than 20% of the total file system space.
- The available space in the file system where the server instance directory is located is less than 1 GB.

If you specify a value of 0, only one log file, db2diag.log, is used for all diagnostic messages. No limits are imposed on the size of the log file.

Restriction: You must monitor the size of the diagnostic log files to ensure that they do not use all the available space in the file system. If there is not enough available space, the server might fail to respond.

Syntax

```
.-1024-----.  
>>-DBDIAGLOGSize--+-megabytes+-----><
```

Parameters

megabytes

Specifies the amount of space that is used by diagnostic log files in megabytes. Specify a value in the range 2 - 9999, or a value of 0. The default value is 1024.

If you specify a value in the range 2 - 9999, rotating log files are used, and the value specifies the total size in megabytes of all 10 log files. The value is reset to 1024 whenever the server is restarted.

If you specify a value of 0, one log file is used, and no limits are imposed on the size of the log file.

If you want to archive messages, specify a value of 0 to ensure that the db2diag.log file can use all the available space without using rotating log files.

After you set the value of the megabytes parameter to 0 by using the DBDIAGLOGSIZE option, messages are initially written to rotating log files. After the server is restarted, messages are written to the db2diag.log file.

Tip: If you specify a value in the range 2 - 9999 by using the server options file, dsmserv.opt, the value is not reset automatically at server startup. The value remains the same until it is changed or removed from the dsmserv.opt file, by using the SETOPT command.

Example: Specify a maximum size of 5120 megabytes

Specify the size of the diagnostic log files as 5120 megabytes (5 GB):

```
dbdiaglogsize 5120
```

Example: Archive messages in a single log file

Archive messages by specifying that the messages are written to the db2diag.log file:

```
dbdiaglogsize 0
```

Related information:

[DB2 V10.5 product information](#)

DBDIAGPATHFSTHRESHOLD

The DBDIAGPATHFSTHRESHOLD option specifies the threshold for free space on the file system or disk that contains the db2diag.log file.

When the amount of free space is equal to or less than the specified threshold, the ANR1545W error message is shown. By default, the message is shown when the file system or disk has 20% or less of free disk space.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-DBDIAGPATHFSTHreshold--percent-----<<
```

Parameter

percent

Specifies the percentage of available space in the file system. Valid values are in the range 0 - 100. The default is 20.

Tip: For best results, do not set a low or high value for the percent parameter. A low value might cause the file system to become full before you can correct the issue. A full file system might corrupt the server database. A high value might result in many ANR1545W messages in the server activity log.

Example

Set the threshold value to 10%.

```
setopt DBDIAGPATHFSTH 10
```

DBMEMPERCENT

Use this option to specify the percentage of the virtual address space that is dedicated to the database manager processes.

If applications other than IBM Spectrum Protect™ server are running on the system, ensure that the value allows adequate memory for the other applications.

Syntax

```
>>-DBMEMPERCENT--+-percent+-----><  
                    '-AUTO-----'
```

Parameters

percent

Set a value from 10 to 99.

AUTO

The database manager sets the percentage automatically to a value that is between 75 percent and 95 percent of system RAM. The default value is AUTO.

Examples

```
dbmempercent 50
```

DBMTCPPORT

The DBMTCPPORT option specifies the port number on which the TCP/IP communication driver for the database manager waits for requests for client sessions.

The specified port number must be reserved for use by the database manager.

By default, the IBM Spectrum Protect™ server uses interprocess communications (IPC) to establish connections for the first two connection pools, with a maximum of 480 connections for each pool. After the first 960 connections are established, the IBM Spectrum Protect server uses TCP/IP for any additional connections.

Syntax

```
>>-DBMTCPPort--port_number-----><
```

Parameters

port_number

Specifies the number of the TCP/IP port on which the database manager waits for communications from the server. Valid values are integers from 1024 to 65535.

The default port number is the value of the server TCPPOINT option plus 50,000. For example, if the server TCPPOINT option is 1500, the default DBMTCPPORT port number would be 51500.

If the TCPPOINT server option is greater than 9999, add the last four digits of its value to 50000. For example, if the TCPPOINT option is 11500, 1550 is added to 50000, resulting in a DBMTCPPORT port number of 51500.

Example

```
dbmtcport 51500
```

DEDUPREQUIRESBACKUP

The DEDUPREQUIRESBACKUP option specifies whether volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and whether duplicate data can be discarded before the storage pools are backed up.

If the value of this option is YES (the default), you must back up data to copy storage pools that are not set up for data deduplication. Use the BACKUP STGPOOL command to back up data to copy storage pools.

Be aware that reclamation of a volume in a storage pool that is set up for data deduplication might not occur when the volume first becomes eligible. The server makes additional checks to ensure that data from a storage pool that is set up for data deduplication has been backed up to a copy storage pool. These checks require more than one BACKUP STGPOOL instance before the server reclaims a volume. After the server verifies that the data was backed up, the volume is reclaimed.

You can change this option dynamically using the SETOPT command.

Attention: To minimize the possibility of data loss, do not change the default setting for this server option. Specify a value of NO only if you do not have any copy storage pools and are not performing storage pool backups.

Syntax

```
>>-DEDUPREQUIRESBACKUP---+-Yes-+-----><
      '-No--'
```

Parameters

Yes

Specifies that the storage pool must be backed up before volumes can be reclaimed and before duplicate data can be discarded. This is the default.

No

Specifies that volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and duplicate data can be discarded if the storage pools are not backed up.

Examples

Specify that primary sequential-access storage pools that are set up for data deduplication do not have to be backed up.

```
deduprequiresbackup no
```

DEDUPTIER2FILESIZE

The DEDUPTIER2FILESIZE option specifies at what file size IBM Spectrum Protect™ begins to use Tier 2 data deduplication.

Syntax

```
>>-DEDUPTIER2FILESIZE---nnn-----><
```

Parameters

nnn

Specifies the file size, in gigabytes, at which point the IBM Spectrum Protect server begins to use Tier 2 processing for data deduplication. You can specify a value 20 - 9999. The default is 100.

Note: If the value specified or defaulted to for this option is greater than the value for the SERVERDEDUPTXNLIMIT option, then this option is ignored for server data deduplication. If the value specified or defaulted to for this option is greater than the value for CLIENTDEDUPTXNLIMIT, then this option is ignored for client data deduplication.

Examples

```
deduptier2filesize 550
```

DEDUPTIER3FILESIZE

The DEDUPTIER3FILESIZE option specifies at what file size IBM Spectrum Protect™ begins to use Tier 3 data deduplication.

Syntax

```
>>-DEDUPTIER3FILESIZE--nnn-----><
```

Parameters

nnn

Specifies the file size, in gigabytes, at which point the IBM Spectrum Protect server begins to use Tier 3 processing for data deduplication. You can specify a value 90 - 9999. The default is 400.

- If the value specified or defaulted to for this option is greater than the value for the SERVERDEDUPTXNLIMIT option, then this option is ignored for server data deduplication.
- If the value specified or defaulted to for this option is greater than the value for CLIENTDEDUPTXNLIMIT, then this option is ignored for client data deduplication.
- If the value specified or defaulted to for this option is less than the value specified or defaulted to for DEDUPTIER2FILESIZE, then the value of DEDUPTIER2FILESIZE is used for this option.

Examples

```
deduptier3filesize 1150
```

DEVCONFIG

The DEVCONFIG option specifies the name of a file in which you want IBM Spectrum Protect™ to store a backup copy of device configuration information.

IBM Spectrum Protect stores the following information in the device configuration file:

- Device class definitions created by using the DEFINE DEVCLASS command
- Drive definitions created by using the DEFINE DRIVE command
- Library definitions created by using the DEFINE LIBRARY command
- Library inventory information for the LIBTYPE=SCSI automated libraries
- Path definitions created by using the DEFINE PATH command
- Server definitions created with the DEFINE SERVER command
- Server name created with the SET SERVERNAME command
- Server password created with the SET SERVERPASSWORD command

Note:

- Only path definitions with SRCTYPE=SERVER are backed up to the device configuration file. Paths of SRCTYPE=DATAMOVER are not written to the file.
- Library volume location information is stored as comments (*/*...*/*) in the device configuration file whenever CHECKIN LIBVOLUME, CHECKOUT LIBVOLUME, and AUDIT LIBRARY commands are issued for SCSI libraries.

Attention: To restore the database after a disaster, you must have a copy of the current device configuration file. The device configuration file cannot be recreated.

You can include one or more DEVCONFIG options in the server options file. When you use multiple DEVCONFIG options, IBM Spectrum Protect automatically updates and stores a backup copy of device configuration information in each file you specify.

Syntax

```
>>-DEVCONFig--file_name-----><
```

Parameters

file_name

Specifies the name of a file in which to store a backup copy of device configuration information.

Examples

```
devconfig devices.sav
```

DISABLEREORGTABLE

The DISABLEREORGTABLE option specifies whether online table reorganization is disabled for table names that are specified in the tables list.

To use the DISABLEREORGTABLE option, you must halt the server, update the options file, and then restart the server.

Syntax

```
>>-DISABLEREORGTTable----tablelist-----<<
```

Parameters

tablelist

Specifies a list of table names for which table reorganization is disabled. If you do not specify any table names with the option, or if the option is not in the options file, no tables are disabled.

Restriction: The following tables are already excluded from table reorganization processing and cannot be specified for this option:

- STAGED_EXPIRING_OBJECTS
- STAGED_OBJECT_IDS
- BF_DEREFERENCED_CHUNKS
- BF_QUEUED_CHUNKS

Example

```
DISABLEREORGTABLE BF_BITFILE_EXTENTS,REPLICATING_OBJECTS
```

DISABLESCHEDS

The DISABLESCHEDS option specifies whether administrative and client schedules are disabled during IBM Spectrum Protect™ server recovery.

Syntax

```
>>-DISABLESCheds---Yes-----<<  
                '-No--'
```

Parameters

Yes

Specifies that administrative and client schedules are disabled.

No

Specifies that administrative and client schedules are enabled.

Examples

```
disablescheds no
```

DISPLAYLFINFO

The DISPLAYLFINFO option specifies how the accounting records and summary table entries report the node name.

When this option is enabled, the accounting records and summary table entries report node_name(storage_agent_name) for the node name. If the option is not enabled, the accounting records and summary table entries simply report node_name for the node name. The default is No.

Syntax

```
>>-DISPLAYLFINFO--+-Yes+-----><
      '-No--'
```

Parameters

Yes

Specifies that the accounting records and summary table entries will report the storage agent name.

No

Specifies that the accounting records and summary table entries will not report the storage agent name. This is the default.

Examples

```
displaylfinfo yes
```

The result shows the following accounting record with the storage agent name displayed (STA53):

```
5,0,ADSM,07/13/2004,15:35:14,COLIND-TUC (STA53),,WinNT,1,Tcp/Ip,1,0,0,0,
0,223,4063,0,0,222,7,8,3,1,4,0,0,0,0,3,0
```

The corresponding summary table also displays the storage agent name:

```
START_TIME: 2004-07-13 15:35:07.000000
END_TIME: 2004-07-13 15:35:14.000000
ACTIVITY: BACKUP
NUMBER: 8
ENTITY: COLIND-TUC (STA53)
COMMMETH: Tcp/Ip
ADDRESS: colind-tuc:2229
SCHEDULE_NAME:
EXAMINED: 0
AFFECTED: 223
FAILED: 0
BYTES: 4160875
IDLE: 8
MEDIAS: 1
PROCESSES: 1
SUCCESSFUL: YES
VOLUME_NAME:
DRIVE_NAME:
LIBRARY_NAME:
LAST_USE:
COMM_WAIT: 3
NUM_OFFSITE_VOLS:
```

DNSLOOKUP

The DNSLOOKUP option specifies whether the server uses system API calls to determine the domain name server (DNS) names of systems that contact the server.

Syntax

```
>>-DNSLOOKUP--+-Yes+-----><
      '-No--'
```

Parameters

Yes

Specifies that the server obtains the DNS names of contacting systems. Yes is the default.

No

Specifies that the server does not obtain the DNS names of contacting systems.

Examples

```
dnslookup yes
```

DRIVEACQUIRERETRY

The DRIVEACQUIRERETRY option lets you specify how many times the server retries the acquisition of a drive in an IBM® 349x library. If the library is shared among multiple applications, its drives may appear to be available to the server (through the use of a background polling process) when they are not.

This option is only valid if you specified 3494SHARED YES in the dsmserv.opt file. If you specified DRIVEACQUIRERETRY NEVER, you need to monitor how long jobs have been waiting for drives and how long the server has been polling the drives. You may also need to check the status of these drives in the other IBM Spectrum Protect™ servers. There may be cartridges stuck in the drives, and the other IBM Spectrum Protect servers may have marked the drives as *offline*. If this is the case, you need to mark the drives *offline* in the IBM Spectrum Protect server that is polling the drives. If necessary, also cancel any waiting jobs.

Syntax

```
>>-DRIVEACquireretry--+-Forever-----+-----><
                        +-Never-----+
                        '-number_of_retries-'
```

Parameters

Forever

The acquisition of a drive is retried until one is successfully acquired. This is the default.

Never

The server does not retry the acquisition of a drive and fails the operation.

number_of_retries

Specifies the maximum number of times, from 1 to 9999, that the server retries the acquisition of a drive.

Examples

Specify that the server should attempt no more than 10 times to acquire the drive:

```
driveacquireretry 10
```

ENABLENASDEDUP

The ENABLENASDEDUP server option specifies whether the server deduplicates data that is stored by a network-attached storage (NAS) file server. This option applies only to NetApp file servers.

If the value of this option is NO, the data stored by the file server is skipped during duplicate-identification processing. If the value of this option is YES, the value of the DEDUPLICATE parameter in the storage pool definition must be YES.

Syntax

```
>>-ENABLENASDEDUP--+-No-----><
                    '-Yes-'
```

Parameters

Yes

Specifies that IBM Spectrum Protect™ server deduplicates data stored by a NetApp file server.

No

Specifies that the server does not deduplicate data stored by a NetApp file server.

Example

Specify that the server deduplicates data stored by a NetApp file server.

```
enablenasdedup yes
```

EVENTSERVER

The EVENTSERVER option specifies whether at startup the server should try to contact the event server.

Syntax

```
>>-EVENTSERVer--+-Yes-+-----><  
      '-No--'
```

Parameters

Yes

Specifies that, at startup, the server tries to contact the event server. Contact occurs only if a DEFINE EVENTSERVER command has already been issued. This is the default.

No

Specifies that, at startup, the server does not try to contact the event server.

Examples

```
eventserver yes
```

EXPINTERVAL

The EXPINTERVAL option specifies the interval, in hours, between automatic inventory expiration processes by IBM Spectrum Protect™. Inventory expiration removes client backup and archive file copies from the server as specified by the management classes to which the client files are bound. If expiration is not run periodically, storage pool space is not reclaimed from expired client files, and the server requires more storage space than required by policy.

You can also use the EXPIRE INVENTORY command to start inventory expiration. Expiration can make space available in your storage pools for additional client backup or archive files.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
      .-24----.  
>>-EXPINterval---+hours+-----><
```

Parameters

hours

Specifies the time, in hours, between automatic inventory expiration processes. You can specify from 0 to 336 (14 days). A value of 0 means that expiration must be started with the EXPIRE INVENTORY command. The default is 24.

Examples

```
expinterval 5
```

EXPQUIET

The EXPQUIET option specifies whether IBM Spectrum Protect™ sends detailed messages during expiration processing.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-EXPQUIet---+- --No---+-----><
      '- --Yes-'
```

Parameters

No

Specifies that the server sends detailed messages. This is the default.

Yes

Specifies that the server sends only minimal messages. These messages are sent only for files that have expired based on the copy group in the default management class or retention grace period for the domain.

Examples

```
expquiet no
```

Linux

FASPBEGPORT

The FASPBEGPORT option specifies the starting number in the range of port numbers that are used for network communications with Aspera® Fast Adaptive Secure Protocol (FASP®) technology.

To define the range of port numbers, specify both the FASPBEGPORT and FASPENDPORT options.

Syntax

```
      .-15100-----.
>>-FASPBEGPort---+starting_port_number+-----><
```

Parameters

starting_port_number

Specifies the starting port number for network communications that use Aspera FASP technology. The default value is 15100.

Ask your network administrator to help you define the range of port numbers:

- If you did not enable the Secure Sockets Layer (SSL) protocol for the server pair, ensure that the ports can be used for Transmission Control Protocol (TCP) sockets.
- Ensure that the ports can be used for User Datagram Protocol (UDP) connections.
- Ensure that the ports are compatible with firewall rules.

Example

If firewall rules require port numbers to be greater than 1800, you would specify a minimum port number of 1801:

```
faspbegport 1801
```

Related reference:

FASPENDPORT

Linux

FASPENDPORT

The FASPENDPORT option specifies the ending number in the range of port numbers that are used for network communications with Aspera® Fast Adaptive Secure Protocol (FASP®) technology.

To define the range of port numbers, specify both the FASPBEGPORT and FASPENDPORT options.

Syntax

```
.-15199-----.  
>>-FASPENDPort---+ending_port_number+-----<<
```

Parameters

ending_port_number

Specifies the ending port number for network communications that use Aspera FASP technology. The default value is 15199.

Ask your network administrator to help you define the range of port numbers:

- If you did not enable the Secure Sockets Layer (SSL) protocol for the server pair, ensure that the ports can be used for Transmission Control Protocol (TCP) sockets.
- Ensure that the ports can be used for User Datagram Protocol (UDP) connections.
- Ensure that the ports are compatible with firewall rules.

Example

If firewall rules require port numbers to be less than 1900, you would specify a maximum port number of 1899:

```
faspendport 1899
```

Related reference:

FASPBEGPORT

Linux

FASPTARGETRATE

The FASPTARGETRATE option specifies the target rate for data transfer with Aspera® Fast Adaptive Secure Protocol (FASP®) technology. By specifying the target rate, you limit the bandwidth of each network connection that uses Aspera FASP technology. In this way, you can ensure that sufficient bandwidth is available for all network connections.

Syntax

```
.-250000-----.  
>>-FaspTargetRate---+target_rate+-----<<
```

Parameters

target_rate

Specifies the maximum rate, in kilobits per second, for data transfer during a session. The default value is 250000. You can specify values in the range 100 - 100000000.

For example, if you issue the PROTECT STGPOOL command to run two parallel operations at the default target rate, the aggregated throughput does not exceed 500,000 kbps. If your file system can support two operations to protect storage pools at much higher rates than 500,000 kbps of aggregated throughput, and sufficient network bandwidth is available, you can increase the target rate.

To determine the appropriate target rate, consult your network administrator.

Examples

If the allotted network bandwidth is 150,000 kbps, you can set the target rate to 75,000 and use the default number of sessions (two) for the PROTECT STGPOOL command.

```
fasptargetrate 75000
```

In a large blueprint configuration, if the allotted network bandwidth is 6,000,000 kbps, you can set the target rate to 750,000 and use eight sessions for the PROTECT STGPOOL command.

```
fasptargetrate 750000
```

FFDCLOGLEVEL

The FFDCLOGLEVEL option specifies the type of general server messages that are displayed in the first failure data capture (FFDC) log.

The FFDC log contains three categories of general server messages. Setting the FFDCLOGLEVEL option affects the following categories:

- FFDC_GENERAL_SERVER_INFO
- FFDC_GENERAL_SERVER_WARNING
- FFDC_GENERAL_SERVER_ERROR

Syntax

```
.-FFDCLOGLevel-----ALL-----.  
>>-+-FFDCLOGLevel-----+--ALL--+----->>  
                               +-WARN--+  
                               '-ERRor-'
```

Parameters

ALL

Specifies that all FFDC general server log messages are in the log. This value is the default.

WARN

Specifies that the FFDC_GENERAL_SERVER_WARNING and FFDC_GENERAL_SERVER_ERROR messages appear in the log.

ERRor

Specifies that only the FFDC_GENERAL_SERVER_ERROR messages appear in the log.

Example

```
ffdcloglevel warn
```

FFDCLOGNAME

The FFDCLOGNAME option specifies a name for the first failure data capture (FFDC) log.

The FFDC log file is used to gather diagnostic information about the server. When an error occurs, data about the error is written to the FFDC log file. This information can be provided to IBM Support to help diagnose problems. The FFDC log file is in the server instance directory.

Syntax

```
.-dsmffdc.log-  
>>-FFDCLOGNAME---+file_name----->>
```

Parameters

file_name

Specifies a file name for the FFDC log file. The file name can be a fully qualified file name or a file name relative to the server instance directory. The default value is dsmffdc.log.

Examples

```
ffdclogname /tsminst1/tsmffdc.log
ffdclogname tsmffdc.log
ffdclogname c:\tsmserv1\tsmffdc.log
```

Related reference:

FFDCMAXLOGSIZE
FFDCNUMLOGS

FFDCMAXLOGSIZE

The FFDCMAXLOGSIZE option specifies the size for the first failure data capture (FFDC) log file.

The FFDC log file is used to gather diagnostic information about the server. When an error occurs, data about the error is written to the FFDC log file. This information can be provided to IBM Support to help diagnose problems.

Syntax

```
                .-1024-----.
>>-FFDCMAXLOGSIZE--+-kilobytes+-----><
```

Parameters

kilobytes

Specifies the size to which the FFDC log file can grow before wrapping. The minimum value is 500. The maximum value is 2097151. The default value is 1024.

To allow the size of the log file to grow indefinitely, specify a value of -1. To disable the log, specify 0.

Examples

```
ffdcmaxlogsize 2000
```

Related reference:

FFDCLOGNAME
FFDCNUMLOGS

FFDCNUMLOGS

The FFDCNUMLOGS option specifies the number of log files that can be used for circular logging. The default value is 10.

Circular logging uses a ring of log files to provide recovery from transaction failures and system crashes. For example, when the dsmffdc.log file is full, it is renamed to dsmffdc.log.1. If a dsmffdc.log.1 file exists, the dsmffdc.log.1 file is renamed to dsmffdc.log.2. If a dsmffdc.log.2 exists, the dsmffdc.log.2 file is renamed to dsmffdc.log.3, and so on, until the FFDCNUMLOGS value is reached. If there is a log file that is renamed as the FFDCNUMLOGS value is reached, that log file is deleted.

The minimum value is 1. The maximum value is 100. The default value is 10.

Syntax

```
                .-10-----.
>>-FFDCNUMLOGS--+-value+-----><
```

Parameters

value

Specifies the number of log files that are used for circular logging.

If you specify a value of 1 and the log file size reaches the FFDCMAXLOGSIZE, the server continues to write to the log file. Any logging information is overwritten and the server continues to write to the log file.

Examples

```
ffdcnumlogs 20
```

FILEEXIT

The FILEEXIT option specifies a file to which enabled events are routed. Each logged event is a record in the file.

Syntax

```
>>-FILEEXIT---No---file_name---REPLACE---<-----<
      '-Yes-'                +-APPEND---+
                              '-PRESERVE-'
```

Parameters

Yes

Specifies that event logging to the file exit receiver begins automatically at server startup.

No

Specifies that event logging to the file exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the BEGIN EVENTLOGGING command.

file_name

Specifies the name of the file in which the events are stored.

REPLACE

Specifies that if the file already exists, it will be overwritten.

APPEND

Specifies that if the file already exists, data is appended to it.

PRESERVE

Specifies that if the file already exists, it will not be overwritten.

Examples

Windows

```
fileexit yes \tssm\server\data replace
```

AIX

Linux

```
fileexit yes /tssm/server/data replace
```

FILETEXTTEXT

The FILETEXTTEXT option specifies a file to which enabled events are routed. Each logged event is a fixed-size, readable line.

Syntax

```
>>-FILETEXTTEXT---No---file_name---REPLACE---<-----<
      '-Yes-'                +-APPEND---+
                              '-PRESERVE-'
```

Parameters

Yes

- Specifies that event logging to the file exit receiver begins automatically at server startup.
- No
 - Specifies that event logging to the file exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the BEGIN EVENTLOGGING command.
- file_name
 - Specifies the name of the file in which the events are stored.
- REPLACE
 - Specifies that if the file already exists, it will be overwritten.
- APPEND
 - Specifies that if the file already exists, data will be appended to it.
- PRESERVE
 - Specifies that if the file already exists, it will not be overwritten.

Examples

Windows

```
filetextexit yes \tsm\server\data replace
```

AIX

Linux

```
filetextexit yes /tsm/server/data replace
```

FIPSMODE

The FIPSMODE option specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for non-Secure Sockets Layer (SSL) operations.

Syntax

```
.-FIPSMODE-----No-----
>>+-----+----->>
'-FIPSMODE-----+No--+-'
      '-Yes-'
```

Parameters

- No
 - Specifies that FIPS mode is not enforced on the server for non-SSL operations. The default is NO.
- Yes
 - A value of YES indicates that FIPS mode is enforced on the server. This setting restricts cryptographic operations that involve object data, authentication, and passwords to use FIPS-approved cipher suites. The value does not affect SSL session operations, which are controlled by using the SSLFIPSMODE option.

Example: Enable FIPS mode on the server

```
fipsmode yes
```

Example: Enable FIPS mode and SSLFIPS mode on the server

```
fipsmode yes
sslfipsmode yes
```

FSUSEDTHRESHOLD

The FSUSEDTHRESHOLD option specifies what percentage of the file system can be filled up by the database before an alert message is issued.

You can update this server option without stopping and restarting the server by using the SETOPT command.

If this value is set to a low number, the activity log might be flooded with messages about the database space being filled, even if there is still space available. If the value is set too high, the database space might be filled before you can add more space to the file system.

Syntax

```
>>-FSUSEDThreshold--percent-----><
```

Parameters

percent

Specifies the value of used space in the database. You can specify a value from 0 to 100. The default is 90.

Examples

```
fsusedthreshold 70
```

IDLETIMEOUT

The IDLETIMEOUT option specifies the amount of time, in minutes, that a client session can be idle before the server cancels the session. You may want to increase the time-out value to prevent clients from timing out if there is a heavy network load in your environment. Note, however, that a large number of idle sessions could prevent other users from connecting to the server.

The IDLETIMEOUT server option is used for non-administrative sessions. See the ADMINIDLETIMEOUT option for administrative client sessions.

You can update this server option without stopping and restarting the server by using the SETOPT command.

Syntax

```
.-15-----.  
>>-IDLETimeout--+minutes+-----><
```

Parameters

minutes

Specifies the maximum number of minutes that a server waits for an idle client. The default value is 15 minutes. The minimum value is 1 minute.

Examples

```
idletimeout 15
```

KEEPALIVE

The KEEPALIVE option specifies whether the Transmission Control Protocol (TCP) keepalive function is enabled for outbound TCP sockets. The TCP keepalive function sends a transmission from one device to another to check that the link between the two devices is operating.

If you are using node replication, you can use the KEEPALIVE option on the source replication server to enable the TCP keepalive function. The KEEPALIVE option is not required on the target replication server unless you specify bidirectional replication, in which case the target server becomes the source replication server.

Syntax

```
.-Yes-.
```

```
>>-KEEPALIVE---+No---+-----><
```

Parameters

Yes

Specifies that the TCP keepalive function is enabled for outbound TCP sockets. This value is the default. If the KEEPALIVE option is enabled, default values are used for the KEEPALIVETIME and KEEPALIVEINTERVAL options.

No

Specifies that the TCP keepalive function is not enabled for outbound TCP sockets. If you specify a value of NO, it does not affect current TCP socket connections that originated from outbound connection requests while the KEEPALIVE option was set to YES. The YES value applies to those sockets until the related session ends and the socket is closed.

Example

Use the SETOPT command to enable the keepalive function without disabling or halting the server:

```
setopt keepalive yes
```

Related reference:

KEEPALIVEINTERVAL
KEEPALIVETIME

KEEPALIVETIME

The KEEPALIVETIME option specifies how often TCP sends a keepalive transmission when it receives a response. This option applies only if you set the KEEPALIVE option to YES.

Syntax

```
.-300-----.  
>>-KEEPALIVETIME---+seconds-+-----><
```

Parameters

seconds

Specifies how often TCP sends keepalive transmissions to verify that an idle connection is still active. The value is specified in seconds.

You can specify a value in the range 1 - 4294967. The default is 300 (5 minutes).

Example

Set the KEEPALIVETIME option to 120 seconds:

```
keepalivetime 120
```

Related reference:

KEEPALIVE
KEEPALIVEINTERVAL

KEEPALIVEINTERVAL

The KEEPALIVEINTERVAL option specifies how often a keepalive transmission is sent if no response is received. This option applies only if you set the KEEPALIVE option to YES.

Syntax

```
.-30-----.
```

```
>>-KEEPALIVEINTERVAL--+-seconds+-----><
```

Parameters

seconds

Specifies the length of time, in seconds, between keepalive transmissions when no response is received. The value is specified in seconds.

You can specify a value in the range 1 - 4294967. The default is 30 seconds.

Example

Set the KEEPALIVEINTERVAL option to 45 seconds:

```
keepaliveinterval 45
```

Related reference:

KEEPALIVE

KEEPALIVETIME

LANGUAGE

The LANGUAGE option controls the initialization of locales. A locale includes the language and the date, time, and number formats to be used for the console and server.

If your client and server are running different languages, the messages that are generated might not be understandable when messages are issued from the client to the server or if the server sends output to the client.

AIX | **Linux** If initialization of the locale fails, the server defaults to American English.

Windows If the initialization of the locale fails, the server defaults to American English and uses the date, time, and number formats that are set by the DATEFORMAT, TIMEFORMAT, and NUMBERFORMAT server options.

Syntax

```
>>-LANGUage--+-AMENG---(1)-----><
      |          (2) |
+-en_US-----+
      |          (3) |
'-locale-----'
```

Notes:

1. AMENG is available only on HP-UX, Solaris, Windows.
2. en_US is available only on AIX and Linux.
3. *locale* is available only on AIX, HP-UX, Solaris, Linux, and Windows.

Parameters

Windows AMENG

Windows Specifies that American English is used as the default language for the server.

AIX | **Linux** en_US

AIX | **Linux** Specifies that American English is used as the default language for the server.

locale

Specifies the name of the locale that is supported by the server. See the following tables for information on supported locales by operating system.

Note: IBM Spectrum Protect™ runs in any locale, but defaults to American English. For the locales listed, language support is available.

AIX

Table 1. Server languages for AIX®

Language	LANGUAGE option value
Chinese, Simplified	zh_CN
Chinese, Simplified	Zh_CN
Chinese, Simplified (UTF-8)	ZH_CN
Chinese, Traditional (Big5)	Zh_TW
Chinese, Traditional (UTF-8)	ZH_TW
Chinese, Traditional (euc_tw)	zh_TW
English	en_US
English (UTF-8)	EN_US
French	fr_FR
French (UTF-8)	FR_FR
German	de_DE
German (UTF-8)	DE_DE
Italian	it_IT
Italian (UTF-8)	IT_IT
Japanese, EUC	ja_JP
Japanese, PC	Ja_JP
Japanese, UTF8	JA_JP
Korean	ko_KR
Korean (UTF-8)	KO_KR
Portuguese, Brazilian	pt_BR
Portuguese, Brazilian (UTF-8)	PT_BR
Russian	ru_RU
Russian (UTF-8)	RU_RU
Spanish	es_ES
Spanish (UTF-8)	ES_ES
Table note: The system must have en_US environment support installed.	

Linux

Table 2. Server languages for Linux

LANGUAGE	LANGUAGE option value
Chinese, Simplified	zh_CN
	zh_CN.gb18030
	zh_CN.utf8
Chinese, Traditional	Big5 / Zh_TW
	zh_TW
	zh_TW.utf8
English, United States	en_US
	en_US.utf8
French	fr_FR
	fr_FR.utf8
German	de_DE
	de_DE.utf8

LANGUAGE	LANGUAGE option value
Italian	it_IT
	it_IT.utf8
Japanese	ja_JP
	ja_JP.utf8
Korean	ko_KR
	ko_KR.utf8
Portuguese, Brazilian	pt_BR
	pt_BR.utf8
Russian	ru_RU
	ru_RU.utf8
Spanish	es_ES
	es_ES.utf8

Windows

Table 3. Server languages for Windows

Language	LANGUAGE option value
Chinese, Simplified	chs
Chinese, Traditional	cht
English	ameng
French	fra
German	deu
Italian	ita
Japanese	jpn
Korean	kor
Portuguese, Brazilian	ptb
Russian	rus
Spanish	esp

Examples

AIX Linux

```
lang ja_JP
```

Windows

```
lang jpn
```

LDAPCACHEDURATION

The LDAPCACHEDURATION option determines the amount of time that the IBM Spectrum Protect™ server caches LDAP password authentication information.

After a successful LDAP bind, the value that you enter determines the amount of time that information about the LDAP directory server is kept available. The higher the number, the better the performance of the LDAP directory server. During the cache period, though, changes on the LDAP directory server do not take immediate effect on the node. For example, old passwords might be available for some time, even after they were changed or locked on the LDAP server.

Include the LDAPCACHEDURATION option in a SETOPT command to have the option take effect immediately.

Restriction: The LDAPCACHEDURATION option does not apply to storage agents.

Syntax

```
>>-LDAPCACHEDURATION--minutes-----<<
```

Parameters

minutes

Specifies the maximum amount of time after a successful LDAP bind, that subsequent sessions to the same node or administrator skip secondary LDAP bind operations. Values range from zero to 360 minutes.

Example: Set the LDAPCACHEDURATION value to 6 hours (maximum)

In the dsmserv.opt file, specify the following value:

```
ldapcacheduration 360
```

After a node or administrator authenticates with an external directory server, the LDAP bind is skipped for 360 minutes on all sessions.

LDAPURL

The LDAPURL option specifies the location of a Lightweight Directory Access Protocol (LDAP) server. Set the LDAPURL option after you configure the LDAP server.

Tip: The information in this documentation applies to the LDAP authentication method that is preferred for IBM Spectrum Protect™ V7.1.7 or later servers. For instructions about using the previous LDAP authentication method, see [Managing passwords and logon procedures](#).

The following restrictions apply:

- The LDAPURL option cannot be used in combination with the SETOPT command.
- The LDAPURL option does not apply to storage agents.

Syntax

```
>>-LDAPURL--ldap_url_value-----<<
```

Parameters

ldap_url_value

Specifies the URL of one LDAP server, or the URLs of multiple LDAP servers. You can enter multiple values, with each URL value up to 1024 characters. The port number is optional and defaults to 389. Each URL value must contain an LDAP server name. For example, the format of the server name is `server1.storage.us.ibm.com` and the LDAP port is 341. The value of the LDAPURL option must conform to the following specifications:

- If you specify multiple URLs, each URL must be on a separate line.
- If you specify multiple URLs, each URL must point to a different external directory, and all external directories must contain the same data.
- Each URL must begin with `ldap://`.

Restriction: The URL that you designate cannot begin with `ldaps://`.

IBM Spectrum Protect supports LDAP connections that are secured with the standard LDAPv3 StartTLS operation, which establishes a secure Transport Layer Security (TLS) exchange on an existing LDAP connection. The LDAP Simple Bind operation that IBM Spectrum Protect uses does not protect the password when it is sent. A secure TLS connection is required to protect the password.

Example: Set the port value for an LDAP server

In the dsmserv.opt file, specify the port value as 341 for an LDAP server:

MAXSESSIONS

The MAXSESSIONS option specifies the maximum number of simultaneous client sessions that can connect with the server.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
                .-25-----.  
>>-MAXSessions--+-number_of_sessions-+-----<<
```

Parameters

number_of_sessions

Specifies the maximum number of simultaneous client sessions. The default value is 25 client sessions. The minimum value is 2 client sessions. The maximum value is limited only by available virtual storage size or communication resources.

Examples

```
maxsessions 25
```

MESSAGEFORMAT

The MESSAGEFORMAT option specifies whether a message number is displayed in all lines of a multi-line message.

Syntax

```
>>-MMessageformat--number-----<<
```

Parameters

number

Select a number to specify if a message number is to be displayed only on the first line of a multi-line message or is to be displayed on all lines.

1

The message number for a message is displayed only in the first line of the message. This is the default.

2

The message number for a message is displayed in all lines of a message.

Examples

```
messageformat 2
```

MIRRORLOGDIRECTORY

The MIRRORLOGDIRECTORY option specifies the directory for mirroring the active log path.

All changes made to the active log directory are also written to this mirror directory. This option is appended to the options file when the DSMSEV FORMAT command is run. Typically, the directory does not need to be changed.

Syntax

```
>>-MIRRorlogdirectory--dir_name-----<<
```

Parameters

dir_name

Specifies a fully qualified directory name for the active log mirror. The maximum number of characters is 175.

Examples

AIX Linux

```
mirrorlogdirectory /tsm/mirrorlog
```

Windows

```
mirrorlogdirectory c:\tsmserv1\mirrorlog
```

MOVEBATCHSIZE

The MOVEBATCHSIZE option specifies the number of client files that are to be moved and grouped together in a batch, within the same server transaction. This data movement results from storage pool backups and restores, migration, reclamation, and MOVE DATA operations. This option works with the MOVESIZETHRESH option.

Syntax

```
                .-1000-----.  
>>-MOVEBatchsize--+-number_of_files-+-----<<
```

Parameters

number_of_files

Specifies a number of files between 1 and 1000. The default is 1000.

Examples

```
movebatchsize 100
```

MOVESIZETHRESH

The MOVESIZETHRESH option specifies, in megabytes, a threshold for the amount of data moved as a batch, within the same server transaction. When this threshold is reached, no more files are added to the current batch, and a new transaction is started after the current batch is moved.

Syntax

```
                .-4096-----.  
>>-MOVESizethresh--+- megabytes-+-----<<
```

Parameters

megabytes

Specifies the number of megabytes as an integer from 1 to 32768. The default value is 4096. This option is used with the MOVEBATCHSIZE option.

Examples

```
movesizethresh 500
```

MSGINTERVAL

The MSGINTERVAL option specifies the time, in minutes, between messages prompting an operator to mount a tape for the server.

Syntax

```
                .-1-----.  
>>-MSGINTERval--+-minutes+-----><
```

Parameters

minutes

Specifies the time interval at which the operator is prompted by the server to mount a tape. The default value is 1 minute. The minimum value is 1 minute.

Examples

```
msginterval 2
```

Windows

NAMEDPIPENAME

The NAMEDPIPENAME option specifies a communication method that allows processes to communicate with one another without having to know where the sender and receiver processes are located. The name acts like an alias, connecting the two processes regardless of whether they are on the same computer or across connected domains.

Syntax

```
>>-NAMEDpipename--name-----><
```

Parameters

name

Specifies the named pipes name for the server to use. Named pipes are ideal for running in an environment where client and server are on the same machine. No communication software is required and no setup is required.

Examples

```
namedpipename  \\.\PIPE\TSMPIPE
```

AIX

Linux

Windows

NDMPCONNECTIONTIMEOUT

The NDMPCONNECTIONTIMEOUT server option specifies the time in hours that IBM Spectrum Protect™ server waits to receive status updates during NDMP restore operations across the LAN. NDMP restore operations of large NAS file systems can have long periods of inactivity. The default is 6 hours.

Syntax

```
                .-6-----.  
>>-NDMPCONNECTIONTIMEOUT--+-hours+-----><
```

Parameters

hours

The number of hours that the IBM Spectrum Protect server waits to receive status updates during an NDMP restore operation over the LAN. The default value is 6. The minimum is 1 hour. The maximum is 48 hours.

Example

Specify a timeout of 10 hours before the NDMP connection times out:

```
ndmpconnectiontimeout 10
```

NDMPCONTROLPORT

The NDMPCONTROLPORT option specifies the port number to be used for internal communications for certain Network Data Management Protocol (NDMP) operations. The IBM Spectrum Protect™ server does not function as a general purpose NDMP tape server.

Syntax

```
.-10000-----.  
>>-NDMPControlport---+port_number+-----<<
```

Parameters

port_number

The port number to be used for internal communications for certain NDMP operations. The port number must be from 1024 to 32767. The default is 10000.

Examples

```
ndmpcontrolport 9999
```

NDMPENABLEKEEPALIVE

The NDMPENABLEKEEPALIVE server option specifies whether the IBM Spectrum Protect™ server enables Transmission Control Protocol (TCP) keepalive on network data-management protocol (NDMP) control connections to network-attached storage (NAS) devices. The default is NO.

TCP keepalive is implemented within the network support of an operating system. TCP keepalive prevents a long-running, inactive connection from being closed by firewall software that detects and closes inactive connections.

Restriction: To prevent errors, do not enable TCP keepalive in certain types of environments. One example is environments that do not have firewalls between the IBM Spectrum Protect server and a NAS device. Another example is environments with firewalls that tolerate long-running, inactive connections. Enabling TCP keepalive in this type of environment can cause an idle connection to be inadvertently closed if the connection partner temporarily fails to respond to TCP keepalive packets.

Syntax

```
>>-NDMPENABLEKEEPALIVES---+NO---+-----<<  
                                  '-YES-'
```

Parameters

NO

Disable TCP keepalive on all NDMP control connections. NO is the default.

YES

Enable TCP keepalive on all NDMP control connections. The default idle time before the first TCP keepalive packet is sent is 120 minutes.

AIX | **Linux** | **Windows** To change the idle time, use the NDMPKEEPIDLEMINUTES server option.

Example

Enable TCP keepalive on all NDMP control connections so that inactive NDMP connections are not closed:

```
ndmpenablekeepalive yes
```

AIX

Linux

Windows

NDMPKEEPIDLEMINUTES

The NDMPKEEPIDLEMINUTES server option specifies the amount of time, in minutes, before the operating system transmits the first Transmission Control Protocol (TCP) keepalive packet on a network data-management protocol (NDMP) control connection. The default is 120 minutes.

Prerequisite: Use this option only after you set the value of the NDMPENABLEKEEPALIVES server option to YES.

Syntax

```
                .-120-----.  
>>-NDMPKEEPIDLEMINUTES--+-minutes-+-----<<
```

Parameters

minutes

The number of minutes of inactivity on NDMP control connections before TCP keepalive packets are transmitted. The default value is 120. The minimum is 1 minute. The maximum is 600 minutes.

Example

Specify an idle time of 15 minutes before the first TCP keepalive packet is sent:

```
ndmpkeepidleminutes 15
```

NDMPPORTRANGE

The NDMPPORTRANGE option specifies the range of port numbers through which IBM Spectrum Protect™ cycles to obtain a port number for accepting a session from a network-attached storage (NAS) device for data transfer. The default is 0,0 which means that IBM Spectrum Protect lets the operating system provide a port (ephemeral port).

If all ports specified are in use when a NAS device attempts to connect to the server, the operation fails. If a single port number is chosen (no comma and no port number for the high value), the default for the high port number is the low port number plus 100.

When Network Data Management Protocol (NDMP) data is directed to an IBM Spectrum Protect native pool, communication can be initiated from either the NDMP systems or the IBM Spectrum Protect server. If a firewall separates the server and NAS devices, it may be necessary to specify port numbers in firewall rules to allow traffic to pass to and from the NAS devices. NAS devices communicate to the IBM Spectrum Protect server the port numbers that they will use when contacting the server. The port numbers of the server are controlled with the NDMPPortrange options. Port number control for NAS devices is specific to vendors. Consult your vendor documentation.

Syntax

```
>>-NDMPPortrange--port_number_low+-----+-----<<  
                ',port_number_high'
```

Parameters

port_number_low

The low port number from which IBM Spectrum Protect starts to cycle when needing a port number for accepting session from a NAS device for data transfer. The minimum port number value is 1024.

port_number_high

The high port number to which IBM Spectrum Protect can cycle when needing a port number for accepting session from a NAS device for data transfer. The maximum port number value is 32767. The high port number must be the same or larger than the low port number.

Examples

Specify that IBM Spectrum Protect can cycle from port numbers 1024 - 2024.

```
ndmpportrange 1024,2024
```

NDMPREFDATAINTERFACE

This option specifies the IP address that is associated with the interface in which you want the server to receive all Network Data Management Protocol (NDMP) backup data.

This option affects all subsequent NDMP filer-to-server operations, but does not affect NDMP control connections, which use the system's default network interface. The value for this option is a host name or IPV4 address that is associated with one of the active network interfaces of the system on which the IBM Spectrum Protect™ server is running. This interface must be IPV4 enabled.

You can update this server option without stopping and restarting the server by using the SETOPT command.

Syntax

```
>>-NDMPREFDATAINTERFACE--ip_address-----<<
```

Parameters

ip_address

Specify an address in either dotted decimal or host name format. If you specify a dotted decimal address, it is not verified with a domain name server. If the address is not correct, it can cause failures when the server attempts to open a socket at the start of an NDMP filer-to-server backup.

Host name format addresses are verified with a domain name server. There is no default value. If a value is not set, all NDMP operations use the IBM Spectrum Protect server's network interface for receiving backup data during NDMP filer-to-server backup operations.

To clear the option value, specify the SETOPT command with a null value, "".

Examples:

```
ndmprefdatainterface net1.tucson.ibm.com
```

```
ndmprefdatainterface 9.11.152.89
```

NOPREEMPT

The server allows certain operations to preempt other operations for access to volumes and devices. You can specify the NOPREEMPT option to disable preemption. When preemption is disabled, no operation can preempt another for access to a volume, and only a database backup operation can preempt another operation for access to a device.

For example, a client data restore operation preempts a client data backup for use of a specific device or access to a specific volume.

Syntax

```
>>-NOPREEMPT-----<<
```

Parameters

None

Examples

Disable preemption among server operations:

```
nopreempt
```

NORETRIEVEDATE

The NORETRIEVEDATE option specifies that the server does not update the retrieve date of a file in a disk storage pool when a client restores or retrieves the file. This option and the MIGDELAY storage pool parameter control when the server migrates files.

If you do not specify NORETRIEVEDATE, the server migrates files after they have been in the storage pool for the number of days specified by the MIGDELAY parameter. The number of days is counted from the day that the file was stored in the storage pool or retrieved by a client, whichever is more recent. If you specify NORETRIEVEDATE, the server does not update the retrieve date of a file, and the number of days is counted from the day the file entered the disk storage pool.

If you specify this option and caching is enabled for a disk storage pool, reclamation of cached space is affected. When space is needed in a disk storage pool that contains cached files, the server gets the space by selectively erasing cached copies. Files that have the oldest retrieve dates and occupy the largest amount of space are selected for removal. When you specify NORETRIEVEDATE, the server does not update the retrieve date when a file is retrieved. This may cause cached copies to be removed even though they have recently been retrieved by a client.

Syntax

```
>>-NORETRIEVEDATE-----<<
```

Parameters

None.

Examples

Specify that the retrieve dates of files in disk storage pools are not updated when clients restore and retrieve the files:

```
noretrievedate
```

Windows

NPAUDITFAILURE

The NPAUDITFAILURE option specifies whether an event is sent to the event log when a node logs in to the server using a name that is in the Windows group but does not match the Windows account login name. To ensure that a node can access only its own data, the node name and the Windows account name must match.

Syntax

```
>>-NPAUDITFailure--+Yes+-----<<  
                  '-No--'
```

Parameters

Yes

Specifies that an event is sent to the event log when a node logs in to the server using a name that is in the Windows group. But, this name does not match the Windows account login name.

No
Specifies that an audit failure event is not sent to the event log.

Examples

Specify that an event is sent to the event log when a node logs in to the server using a name that is in the Windows group. But, this name does not match the Windows account login name.

```
npauditfailure yes
```

Windows

NPAUDITSUCCESS

The NPAUDITSUCCESS option specifies that an event is sent to the event log when a client node user is authenticated for access to the server through SECUREPIPE.

Syntax

```
>>-NPAUDITSUCCESS--+-Yes+-----<<  
                    '-No--'
```

Parameters

Yes
Specifies that an event is sent to the event log when a client node user is authenticated for access to the server through SECUREPIPES.

No
Specifies that an event is not sent to the Windows log.

Examples

Specify that an event is sent to the event log when a client node is authenticated for access to the server.

```
npauditsuccess yes
```

Windows

NPBUFFERSIZE

The NPBUFFERSIZE option specifies the size of the Named Pipes communication buffer.

Syntax

```
                .-8-----.  
>>-NPBUFFERSIZE--+-kilobytes+-----<<
```

Parameters

kilobytes
Specifies the size, in kilobytes, of the Named Pipes communication buffer. The default is 8.

Examples

Specify a 16 KB Named Pipes communication buffer:

```
npbuffersize 16
```

Windows

NUMBERFORMAT

The NUMBERFORMAT option specifies the format in which the server displays numbers.

The value of NUMBERFORMAT is overridden by the number formatting definition of the locale if the locale is successfully initialized at server startup. The locale is specified in the LANGUAGE option.

Syntax

```
>>-NUMBERformat--number-----><
```

Parameters

number

Select a number from 1 to 6 to identify the number format used by the server. The default is 1.

1	1,000.00
2	1,000,00
3	1 000,00
4	1 000.00
5	1.000,00
6	1'000,00

Examples

```
numberformat 4
```

NUMOPENVOLSALLOWED

The NUMOPENVOLSALLOWED option specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time.

Input volumes contain data to be read during client-restore operations and server processes, such as reclamation and migration. Use this option to improve performance by reducing the frequency with which volumes are opened and closed.

Each session within a client operation or server process can have as many open FILE volumes as specified by this option. A session is initiated by a client operation or by a server process. Multiple sessions can be started within each.

During a client restore operation, volumes can remain open for the duration of a client restore operation and as long a client session is active. During a no-query restore operation, the volumes remain open until the no-query restore completes. At that time, all volumes are closed and released. However, for a classic restore operation started in interactive mode, the volumes might remain open at the end of the restore operation. The volumes are closed and released when the next classic restore operation is requested.

Set this value in the server options file or use the SETOPT command.

Tip: This option can significantly increase the number of volumes and mount points in use at any one time. To optimize performance, follow these steps:

- To set NUMOPENVOLSALLOWED, select a beginning value (the default is recommended). Monitor client sessions and server processes. Note the highest number of volumes open for a single session or process. Increase the setting of NUMOPENVOLSALLOWED if the highest number of open volumes is equal to the value specified by NUMOPENVOLSALLOWED.
- To prevent sessions or processes from having to wait for a mount point, increase the value of the MOUNTLIMIT parameter in the device-class definition. Set the value of the MOUNTLIMIT parameter high enough to allow all client sessions and

server processes using deduplicated storage pools to open the number of volume specified by the NUMOPENVOLSAALLOWED option. For client sessions, check the destination in the copy group definition to determine how many nodes are storing data in the deduplicated storage pool. For server processes, check the number of processes allowed for each process for the storage pool.

- A situation might occur in which a node backs up and restores or archives and retrieves concurrently to and from a deduplicated storage pool. All the mount points required for these operations increase the total number of mount points required by the node.

As a result, the node might not be able to start additional backup sessions if it already has more mount points open than what the MAXNUMMP parameter in the client-node definition allows. This can occur even though the MOUNTLIMIT for the device class was not exceeded.

To prevent backup and retrieve operations from failing, set the value of the MAXNUMMP parameter in the client-node definition to a value at least as high as the NUMOPENVOLSAALLOWED option. Increase this value if you notice that the node is failing backup or retrieve operations because the MAXNUMMP value is being exceeded.

Syntax

```
>>-NUMOPENVOLSAallowed--number_of_open_volumes-----><
```

Parameters

number_of_open_volumes

Specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time. The default is 10. The minimum value is 3. The maximum value is 999.

Examples

Specify that up to 5 volumes in a deduplicated storage pool can be open at one time.

```
numopenvolsallowed 5
```

AIX | Linux | Windows

PUSHSTATUS

The PUSHSTATUS option is used on spoke servers to ensure that status information is sent to the hub server. Do not update this option unless you must restore the Operations Center configuration to the preconfigured state where the IBM Spectrum Protect™ servers are not defined as hub or spoke servers.

If you must restore the Operations Center configuration to the preconfigured state, you must issue the following command on each spoke server:

```
SETOPT PUSHSTATUS NO
```

QUERYAUTH

The QUERYAUTH option specifies the administrative authority level required to issue QUERY or SQL SELECT commands. By default any administrator can issue QUERY and SELECT commands. You can use this option to restrict the use of these commands.

Syntax

```
>>-QUERYAuth--+-None-----+-----><
      +-System---+
      +-Policy---+
      +-Storage---+
      '-Operator-'
```

Parameters

NOne

Any administrator can issue QUERY or SELECT commands without requiring any administrative authority.

SYstem

Administrators must have SYSTEM authority to issue QUERY or SELECT commands.

POlICY

Administrators must have POLICY authority over one or more policy domains or SYSTEM authority to issue QUERY or SELECT commands.

STorage

Administrators must have STORAGE authority over one or more storage pools or SYSTEM authority to issue QUERY or SELECT commands.

OPerator

Administrators must have OPERATOR or SYSTEM authority to issue QUERY or SELECT commands.

Examples

To restrict the use of QUERY and SELECT commands to administrators with system or storage authority, enter:

```
queryauth storage
```

RECLAIMDELAY

This option delays the reclamation of a SnapLock volume, allowing remaining data to expire so that there is no need to reclaim the volume.

Syntax

```
                .-4-----.  
>>-RECLAIMDELAY--+-number_of_days+-----><
```

Parameters

number_of_days

Specifies the number of days to delay the reclamation of a SnapLock volume.

Before a SnapLock volume is reclaimed, the IBM Spectrum Protect™ server allows the specified number of days to pass, so that any files remaining on the volume have a chance to expire. The default reclaim delay period is 4 days and can be set anywhere from 1 to 120 days.

Examples

Specify that the number of days to delay reclamation is 30 days:

```
reclaimdelay 30
```

RECLAIMPERIOD

This option allows you to set the number of days for the reclamation period of a SnapLock volume.

Syntax

```
                .-30-----.  
>>-RECLAIMPERIOD--+-number_of_days+-----><
```

Parameters

number_of_days

Specifies the number of days that are allowed for the reclamation period of a SnapLock volume.

After the retention of a SnapLock volume has expired, the IBM Spectrum Protect™ server will reclaim the volume within the specified number of days if there is still data remaining on the volume. The default reclaim period is 30 days and can be set

anywhere from 7 to 365 days.

The reclamation period does not begin until the RECLAIMDELAY period has expired.

Examples

Specify that the reclaim period is 45 days:

```
reclaimperiod 45
```

REORGBEGINTIME

The REORGBEGINTIME option specifies the earliest time that the IBM Spectrum Protect™ server can start a table or index reorganization.

Schedule server-initiated reorganizations to start during periods when server activity is low. Use this option together with the REORGDURATION option. The REORGDURATION specifies an interval during which reorganization can start.

Syntax

```
>>-REORGBEGINTime--hh:mm-----><
```

Parameters

hh:mm

Specifies the time that the server can start a reorganization: The default start time 6:00 a.m. Use a 24-hour format to specify the time.

Time	Description	Values
hh	The hour of the day	Specify a number 00 - 23.
mm	The minute of the hour	Specify a number 00 - 59.

Examples

Specify 6:00 a.m. as the earliest time that a reorganization can start.

```
reorgbegintime 06:00
```

Specify 8:30 p.m. as the earliest time that a reorganization can start.

```
reorgbegintime 20:30
```

Specify noon as the earliest time that a reorganization can start.

```
reorgbegintime 12:00
```

Specify 3:30 p.m. as the earliest time that a reorganization can start.

```
reorgbegintime 15:30
```

Specify midnight as the earliest time that a reorganization can start.

```
reorgbegintime 00:00
```

REORGDURATION

The REORGDURATION option specifies an interval during which server-initiated table or index reorganization can start.

Schedule server-initiated reorganizations to start during periods when server activity is low. Use this option together with the REORGBEGINTIME option. The REORGBEGINTIME option specifies the earliest time that the server can start a reorganization.

Syntax

```
>>-REORGDuration--nn-----><
```

Parameters

nn

Specifies the number of hours during which a reorganization can start. The minimum value is 1, the maximum value is 24. The default value is 24.

Example

Specify an interval of four hours during which a reorganization can start.

```
reorgduration 4
```

REPORTRETRIEVE

The REPORTRETRIEVE option reports on restore or retrieve operations that are performed by client nodes or administrators. The default is NO.

Syntax

```
>>-REPORTRETRIEVE---+YES+-----><
      '-NO--'
```

Parameters

YES

Specifies that messages will be issued to the server console and stored in the activity log whenever files are restored or retrieved from the IBM Spectrum Protect™ server. The messages will specify the name of the objects being restored or retrieved and identify the client node or administrator performing the operation.

NO

Specifies that messages will not be issued.

Examples

Specify that messages will be issued and stored in the activity log whenever files are restored or retrieved from the IBM Spectrum Protect server:

```
reportretrieve yes
```

The following message is issued for an administrator client session:

```
ANR0411I Session 8 for administrator COLIND-TUC logged in as node
COLIND-TUC restored or retrieved Backup object: node COLIND-TUC,
filesystem \\colind-tuc\c$, object\CODE\TESTDATA\ XXX.OUT
```

REPLBATCHSIZE

The REPLBATCHSIZE option specifies the number of client files that are to be replicated in a batch, within the same server transaction. This option affects only the node replication processes and works with the REPLSIZETHRESH option to improve node replication processing.

The REPLBATCHSIZE option limits the number of files in a transaction and the REPLSIZETHRESH option limits the number of bytes in a transaction. The transaction ends when either the REPLBATCHSIZE threshold or the REPLSIZETHRESH threshold is reached.

Syntax

```
                .-4096-----.  
>>-REPLBatchsize--+number_of_files+-----><
```

Parameters

number_of_files
Specifies a number of files between 1 - 32768. The default is 4096.

Examples

```
replbatchsize 25000
```

REPLSIZETHRESH

The REPLSIZETHRESH option specifies, in megabytes, a threshold for the amount of data replicated, within the same server transaction.

The amount of data is based on the non-deduplicated size of the file, which is the original size of the file. The amount of data that is replicated is controlled by the threshold. When the amount of data exceeds the threshold, the server ends the transaction and no more files are added to the current batch. A new transaction is started after the current batch is replicated. This option is used with the REPLBATCHSIZE option.

For example, suppose that a file is 10 MB and is stored in a data-deduplication-enabled storage pool and only 2 MB of the file is transferred during replication. The amount of data replicated includes the 10 MB size of the file, and excludes the 2 MB transferred. When the amount of data replicated exceeds the value specified for the REPLSIZETHRESH threshold, the transaction ends.

Tip: If you are replicating data from a source server in the cloud and frequently get an ANR1880W server message on the target server, lower the value of the REPLSIZETHRESH option on the source server.

Syntax

```
                .-4096-----.  
>>-REPLSizethresh--+megabytes+-----><
```

Parameters

megabytes
Specifies the number of megabytes as an integer from 1 - 32768. The default value is 4096.

Examples

```
replsizethresh 2000
```

REQSYSAUTHOUTFILE

The REQSYSAUTHOUTFILE option specifies if system authority is required for administrative commands that cause IBM Spectrum Protect™ to write to an external file.

This option applies to the following commands:

- BACKUP DEVCONFIG with the FILENAMES parameter
- BACKUP VOLHISTORY with the FILENAMES parameter
- DEFINE BACKUPSET
- DELETE BACKUPSET
- GENERATE BACKUPSET
- MOVE DRMEDIA with the CMD parameter
- MOVE MEDIA with the CMD parameter
- QUERY DRMEDIA with the CMD parameter

- QUERY MEDIA with the CMD parameter
- QUERY SCRIPT with the OUTPUTFILE parameter

Syntax

```
>>-REQSYSauthoutfile--+-Yes-+-----><
                        '-No--'
```

Parameters

Yes

System authority is required for administrative commands that cause IBM Spectrum Protect to write to an external file.

No

System authority is not required for administrative commands that cause IBM Spectrum Protect to write to an external file. That is, there is no change to the authority level that is required to issue the command.

Examples

```
reqsysauthoutfile no
```

RESOURCETIMEOUT

The RESOURCETIMEOUT option specifies how long the server waits for a resource before canceling the pending acquisition of a resource. When a timeout occurs the request for the resource will be canceled.

Note: When managing a set of shared library resources, such as servers designated as library managers and clients, consider setting this option at the same time limit for all participants in the shared configuration. In any case of error recovery, IBM Spectrum Protect™ will always defer to the longest time limit.

Syntax

```
                        .-60-----.
>>-RESOURCETimeout--+-minutes-+-----><
```

Parameters

minutes

Specifies the maximum number of minutes that the server waits for a resource. The default value is 60 minutes. The minimum value is 1 minute.

Examples

Specify that the server will wait 15 minutes for a server resource:

```
resourcetimeout 15
```

RESTHTTPSPORT

The RESTHTTPSPORT option specifies the port number to be used for Hypertext Transfer Protocol Secure (HTTPS) communication between the Operations Center and the hub server.

Syntax

```
                        .-8443-----.
>>-RESTHTTPSport--+-secure_port+-----><
```

Parameters

secure_port

Specifies the port number that is used for secure communications between the hub server and the Operations Center. The range of values is 1025 - 32767; the default is 8443.

Example

Specify that port number 8444 is used for HTTPS communication.

```
resthttpsport 8444
```

RESTOREINTERVAL

The RESTOREINTERVAL option specifies how long a restartable restore session can be saved in the server database. As long as the restore session is saved in the database, it can be restarted from the point at which it stopped.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-RESTOREINTERVAL--+-minutes-+-----><
```

Parameters

minutes

Specifies how long, in minutes, that a restartable restore session can be in the database before it can be expired. The minimum value is 0. The maximum is 10080 (one week). The default is 1440 minutes (24 hours). If the value is set to 0 and the restore is interrupted or fails, the restore is still put in the restartable state. However, it is immediately eligible to be expired.

Examples

```
restoreinterval 1440
```

RETENTIONEXTENSION

The RETENTIONEXTENSION option specifies the number of days to extend the retention date of a SnapLock volume. This option allows the server to extend the retention date of a SnapLock volume in order to avoid excessive reclamation.

Syntax

```
>>-RETENTIONEXTENSION--number_of_days-----><
```

Parameters

number_of_days

Specifies the number of days to extend the retention date of a SnapLock volume. The minimum value is 30 days; the maximum value is 9999 days; the default is 365.

If you specify a value of 0 (zero) for the RETVER parameter of an archive copy group, the actual value that is used for RETVER is the value of the option RETENTIONEXTENSION, if one of the following conditions is also true:

- The destination storage pool for the archive copy group is a SnapLock storage pool.
- The storage pool that is the target for a storage pool migration or of a MOVE DATA or MOVE NODEDATA command is a SnapLock storage pool.

If a SnapLock volume is the target volume for data from another SnapLock volume and if the remaining retention of the data on the volume is less than the value specified, then the retention date is set using the value specified. Otherwise, the remaining retention of the data is used to set the retention of the volume.

If a SnapLock volume has entered the reclamation period but the percentage of reclaimable space of the volume has not exceeded the reclamation threshold of the storage pool or the value specified on the THRESHOLD parameter of a RECLAIM STGPOOL command, then the retention date of the SnapLock volume is extended by the amount specified in the RETENTIONEXTENSION option.

Examples

Specify that the retention date is extended by 60 days:

```
retentionextension 60
```

AIX

Linux

Windows

SANDISCOVERY

The SANDISCOVERY option specifies whether the IBM Spectrum Protect™ SAN discovery function is enabled.

To use SAN discovery, all devices on the SAN must have a unique device serial number. When set to ON, the server completes SAN discovery in the following instances:

- When the device path is changed
- When the QUERY SAN command is issued

Using SAN discovery, the server can automatically correct the special file name for a device if it is changed for a specified tape device.

The IBM Spectrum Protect server does not require persistent binding with the SAN discovery function enabled. To display a list of devices that are seen by the server, you can issue the QUERY SAN command.

Syntax

```
.-SANDISCOVERY-----OFF-----  
>>-+-----+----->>  
'-SANDISCOVERY-----+ON-----+'  
          '-UNSCANNEDPATHOFF-'
```

Parameters

ON

Specifies that the server completes SAN discovery when the device path is changed, or when the QUERY SAN command is issued.

OFF

Specifies that the server does not complete SAN discovery when the device path is changed, or when the QUERY SAN command is issued. If the IBM Spectrum Protect server is not able to open a device, a message is issued but the path that is associated with the device is not taken offline. This value is the default.

UNSCANNEDPATHOFF

Specifies that the server does not complete SAN discovery when the device path is changed, or when the QUERY SAN command is issued. If the IBM Spectrum Protect server is not able to open a device, a message is issued and the path to the device is taken offline.

Examples

```
sandiscovery on
```

Related commands

Table 1. Commands related to SANDISCOVERY

Command	Description
---------	-------------

Command	Description
PERFORM LIBACTION	Defines all drives and paths for a library.

AIX Linux Windows

SANDISCOVERYTIMEOUT

The SANDISCOVERYTIMEOUT option specifies the amount of time allowed for host bus adapters to respond when they are queried by the SAN discovery process. Once the time specified for the SANDISCOVERYTIMEOUT is reached, the process times out.

Syntax

```
>>-SANDISCOVERYTIMEOUT--value-----<<
```

Parameters

value

Specifies the amount of time to elapse before the SAN discovery process times out. The range is from 15 to 1800 seconds. The default is 15 seconds.

Examples

```
sandiscoverytimeout 45
```

AIX Linux Windows

SANREFRESHTIME

The SANREFRESHTIME option specifies the amount of time that elapses before the cached SAN discovery information is refreshed. The SANREFRESHTIME option has a default value of 0, which means that there is no SAN discovery cache. The information is obtained directly from the host bus adapter (HBA) every time the server performs a SAN discovery operation.

Note: The QUERY SAN server command always receives SAN information at the time that the command is issued and ignores any value specified for SANREFRESHTIME.

Syntax

```
.-0----.
>>-SANREFRESHTIME--+-time+-----<<
```

Parameters

time

The length of time, in seconds, before the cached SAN discovery information is refreshed. The default value is 0 and specifies that SAN discovery information is not cached. If a value other than 0 is specified, for example, 100 seconds, then the SAN discovery information is refreshed 100 seconds after the prior SAN discovery operation.

Examples

Refresh SAN discovery information after 100 seconds.

```
sanrefreshtime 100
```

Turn off the caching of SAN discovery information.

```
sanrefreshtime 0
```

SEARCHMPQUEUE

The SEARCHMPQUEUE option specifies the order in which the server satisfies requests in the mount queue. If the option is specified, the server first tries to satisfy requests for volumes that are already mounted. These requests may be satisfied before other requests, even if the others have been waiting longer for the mount point. If this option is not specified, the server satisfies requests in the order in which they are received.

Syntax

```
>>-SEARCHMPQUEUE-----<<
```

Parameters

None

Examples

Specify that the server tries to first satisfy a request for a volume that is already mounted:

```
searchmpqueue
```

Windows

SECUREPIPES

When using the named pipes protocol, enabling SECUREPIPES forces the server to check the Windows group designated by ADSMGROUPNAME in order to authenticate a client node/user.

The user name and password defined in the Windows group are used to authenticate the node/user for access to the server data. The node/user must also be a registered IBM Spectrum Protect™ client node. However, the IBM Spectrum Protect client node password is ignored, and the Windows password associated with the user is used.

Syntax

```
>>-SECUREPipes--+Yes+-----<<  
      '-No--'
```

Parameters

Yes

Specifies that IBM Spectrum Protect checks the Windows group designated by ADSMGROUPNAME in order to authenticate a client node/user.

No

Specifies that IBM Spectrum Protect does not check the Windows group designated by ADSMGROUPNAME in order to authenticate a client node/user.

Examples

Specify that IBM Spectrum Protect checks the Windows group to authenticate client nodes.

```
securepipes yes
```

SERVERDEDUPTXNLIMIT

The SERVERDEDUPTXNLIMIT option specifies the maximum size of objects that can be deduplicated on the server.

When you use duplicate-identification processes (the IDENTIFY DUPLICATES command) for large objects, intensive database activity can result from long-running transactions that are required to update the database. High levels of database activity can

produce following symptoms:

- Reduced throughput for client backup and archive operations
- Resource contention resulting from concurrent server operations
- Excessive recovery log activity

The extent to which these symptoms occur depends on the number and size of objects being processed, the intensity and type of concurrent operations taking place on the IBM Spectrum Protect™ server, and the IBM Spectrum Protect server configuration.

With the SERVERDEDUPTXNLIMIT server option, you can specify a maximum size, in gigabytes, for objects that can be deduplicated on the server. If an object or set of objects in a single transaction exceeds the limit specified by SERVERDEDUPTXNLIMIT, the objects are not deduplicated by the server. You can specify a value 32 - 102400 GB. The default value is 5120 GB.

Increasing the value of this option causes the IBM Spectrum Protect server to search for objects previously deferred whose size falls below the new transaction limit.

Remember: The search for objects previously deferred can take time. Use care when increasing the value of SERVERDEDUPTXNLIMIT. Reducing the value of this option does not cause IBM Spectrum Protect to search for deferred objects.

The appropriate value for this option depends on the IBM Spectrum Protect server configuration and concurrent server activity. You can specify a high value for this option if you minimize resource contention. To minimize resource contention, perform operations, such as backup, archive, duplicate identification, and reclamation, at different times.

To update this server option without stopping and restarting the server, use the SETOPT command.

Syntax

```
                .-5120-----.  
>>-SERVERDEDUPTXNlimit--+-gigabytes-+-----<<
```

Parameters

gigabytes

Specifies the maximum size, in gigabytes, of objects that can be duplicated on the server. You can specify a value 32 - 102400. The default value is 5120.

Examples

Disable server-side deduplication for all objects over 120 GB:

```
serverdeduptxnlimit 120
```

SHMPORT

AIX | **Linux** The SHMPORT option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection. **Windows** The SHMPORT option specifies the port that the server listens on for shared memory connections.

Syntax

```
>>-SHMPort--port_number-----<<
```

Parameters

port_number

Specifies the port number. **AIX** | **Linux** You can specify a value from 1024 to 32767. The default value is 1510.

Windows You can specify a value from 1 to 32767. The default value is 1.

Examples

```
shmport 1580
```

```
shmport 1
```

SHREDDING

The SHREDDING option specifies whether shredding of deleted sensitive data is performed automatically or manually. Shredding applies only to data in storage pools that have been explicitly configured to support shredding.

Syntax

```
>>-SHREDDing---+AUTOMATIC+-----><
      '-MANual-----'
```

Parameters

AUTOMATIC

Specifies that shredding occurs automatically as sensitive data is deleted. Use this option to shred sensitive data as soon as possible after it is deleted. If the SHREDDING option is not specified, this is the default behavior. If there is an I/O error during automatic shredding, an error is reported, and shredding of the current object halts. If the I/O error cannot be corrected, you might need to run shredding manually and use the IOERROR keyword.

MANUAL

Specifies that shredding occurs manually, only when the SHRED DATA command is invoked. Use this option to control when shredding takes place, in order to ensure that it does not interfere with other server activities.

Tip: If you specify manual shredding, run the SHRED DATA command regularly, at least as often as you perform other routine server-maintenance tasks (for example, expiration, reclamation, and so on). Doing so can prevent performance degradation of certain server processes (in particular, migration). For best results, run SHRED DATA after any operation (for example, expiration and migration) that deletes files from a shred pool.

Examples

Specify that IBM Spectrum Protect™ automatically shreds data in a storage pool configured for shredding after that data is deleted:

```
shredding automatic
```

SNMPHEARTBEATINTERVAL

The SNMPHEARTBEATINTERVAL option specifies the interval in minutes between queries of the IBM Spectrum Protect™ server.

Syntax

```
>>-SNMPHEARTBEATINTERVAL--minutes+-----><
      .-5-----.
```

Parameters

minutes

Specifies the heartbeat interval in minutes. Valid values are from 0 to 1440 (one day). The default is 5 minutes.

Examples

```
snmpheartbeatinterval 20
```

SNMPMESSAGECATEGORY

The SNMPMESSAGECATEGORY option specifies the trap types used when messages are forwarded from the server, through the Simple Network Management Protocol (SNMP) subagent, to the SNMP manager.

Syntax

```
>>-SNMPMESSAGECATEGORY--+SEVERITY---+-----><
      '-INDIVIDUAL-'
```

Parameters

SEVERITY

Specifies that there are four trap types based on message severity level:

- 1 Severe
- 2 Error
- 3 Warning
- 4 Information

This is the default.

INDIVIDUAL

Specifies that a separate trap type is used for each message. The numeric part of the message identifier indicates the trap type.

Examples

```
snmpmessagecategory individual
```

SNMPSUBAGENT

The SNMPSUBAGENT option specifies the parameters needed for the IBM Spectrum Protect™ subagent to communicate with the Simple Network Management Protocol (SNMP) daemon. This option is only to configure the SNMP subagent for communicating with the SNMP agent; it is ignored by the server.

Syntax

```
>>-SNMPSUBAGENT--+-----+----->
      '-HOSTname--host_name-'
>--+-----+-----+-----><
      '-COMMunityname--community_name-' '-TIMEOUT--seconds-'
```

Parameters

HOSTname host_name

Specifies the TCP/IP name or number of the host running the SNMP agent that the IBM Spectrum Protect SNMP subagent connects to. This parameter is optional. The default name is *localhost*.

COMMunityname community_name

Specifies the configured community name on the system running the SNMP agent. This parameter is optional. The default name is *public*.

TIMEOUT seconds

Specifies the time, in seconds, in which a request must be received. This parameter is optional. The default value is 600.

Examples

```
snmpsubagent hostname jimbo communityname public timeout 2600
```

SNMPSUBAGENTHOST

The SNMPSUBAGENTHOST option specifies the location of the IBM Spectrum Protect™ Simple Network Management Protocol (SNMP) subagent. The default for this option is 127.0.0.1.

Syntax

```
>>-SNMPSUBAGENTHOST--host_name-----<<
```

Parameters

host_name

Specifies the TCP/IP host name or number on which the IBM Spectrum Protect SNMP subagent is located. The subagent and server must be on the same node.

Examples

```
snmpsubagenthost 9.116.23.450
```

SNMPSUBAGENTPORT

The SNMPSUBAGENTPORT option specifies the port number of the IBM Spectrum Protect™ Simple Network Management Protocol (SNMP) subagent.

Syntax

```
>>-SNMPSUBAGENTPORT--port_number-----<<
```

Parameters

port_number

Specifies the port number of the IBM Spectrum Protect SNMP subagent. Valid values are 1000 - 32767. The default is 1521.

Examples

```
snmpsubagentport 1525
```

SSLFIPSMODE

The SSLFIPSMODE option specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for Secure Sockets Layer (SSL). The default is NO.

Because SSLv3 is not supported by FIPS mode, when you are using SSL with Version 6.1 or V5.5 clients, you must turn off FIPS mode.

Syntax

```
.-SSLFIPSMODE-----No-----.  
>>-+-----+-----<<  
'-SSLFIPSMODE-----+No--+-'  
      '-Yes-'
```

Parameters

No

Specifies that SSL FIPS mode is not active on the server. This setting is required when Backup-Archive Client versions previous to IBM Spectrum Protect™ 6.3 are to connect to the server with SSL.

Yes

A value of YES indicates that SSL FIPS mode is active on the server. This setting restricts SSL session negotiation to use FIPS-approved cipher suites. Specifying YES is suggested when SSL communication is activated and all Backup-Archive Clients are at V6.3 or later.

Example: Enable SSL FIPS mode on the server

```
sslipsmode yes
```

SSLINITTIMEOUT

The SSLINITTIMEOUT option specifies the time, in minutes, that the server waits for a Secure Sockets Layer (SSL) session to complete initialization before the server cancels the session.

When you specify this option, an SSL session is canceled if a client, server, or storage agent is not configured for SSL and tries to start an SSL session. Similarly, an SSL session is canceled if a client SSL session and a server are not configured with the same Transport Layer Security (TLS) version. In these situations, the SSL session might fail to completely initialize. The server cancels the session when the specified timeout is reached.

Syntax

```
                .-2-----.  
>>-SSLINITTIMEout--+-minutes-+-----<<
```

Parameters

minutes

Specifies the maximum number of minutes that a server waits for an SSL session to complete initialization. The default value is 2 minutes. The minimum value is 1 minute.

Example

```
sslinittimeout 1
```

SSLTCPADMINPORT

The SSLTCPADMINPORT option specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions only. The sessions are for the command-line administrative client.

Note: Beginning with IBM Spectrum Protect™ Version 8.1.2 and Tivoli® Storage Manager Version 7.1.8, you are no longer required to use the SSLTCPADMINPORT or SSLTCPADMINPORT option to allow SSL-enabled sessions from the client. The port number that is specified in the TCPADMINPORT or TCPADMINPORT option listens for both TCP/IP and SSL-enabled client sessions.

The following types of sessions do not use the Secure Sockets Layer (SSL) protocol:

- Network Data Management Protocol (NDMP)
- Automated Cartridge System Library Software (ACSL)
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require that different port numbers are set for the SSLTCPADMINPORT and SSLTCPADMINPORT options.

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (SSLTCPADMINPORT and SSLTCPADMINPORT):

- When you specify the server's SSL-only port for the LLADDRESS on the DEFINE SERVER or UPDATE SERVER command, you must also specify the SSL=YES parameter.
- When you specify the server's SSL-only port for the client's TCPPOINT option, you must also specify YES for the SSL client option.

The TCP/IP communications driver must be enabled with COMMMETHOD TCPIP or COMMMETHOD V6TCPIP.

Syntax

```
>>-SSLTCPADMINPort--port_number-----<<
```

Parameters

port_number

Specifies the port number of the server. Valid values are 1024 - 32767. There is no default.

Examples

```
ssltcpadminport 1543
```

SSLTCPPOINT

The SSLTCPPOINT option specifies the Secure Sockets Layer (SSL) port number for SSL-enabled sessions only. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.

Important: Beginning with IBM Spectrum Protect™ Version 8.1.2 and Tivoli® Storage Manager Version 7.1.8, you are no longer required to use the SSLTCPPOINT or SSLTCPADMINPORT option to allow SSL-enabled sessions from the client. The port number that is specified in the TCPPOINT or TCPADMINPORT option listens for both TCP/IP and SSL-enabled client sessions.

The following types of sessions do not use SSL:

- Network Data Management Protocol (NDMP)
- Automated Cartridge System Library Software (ACSL)
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require that different port numbers are specified for the SSLTCPADMINPORT and SSLTCPPOINT options.

If you specify the same port number for the SSLTCPPOINT and TCPPOINT options, only SSL connections are accepted and TCP/IP connections are disabled for the port.

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (SSLTCPPOINT and SSLTCPADMINPORT):

- When you specify the server's SSL-only port for the LLADDRESS on the DEFINE SERVER or UPDATE SERVER command, you must also specify the SSL=YES parameter.
- When you specify the server's SSL-only port for the client's TCPPOINT option, you must also specify YES for the SSL client option.

The TCP/IP communications driver must be enabled with COMMMETHOD TCPIP or COMMMETHOD V6TCPIP.

Syntax

```
>>-SSLTCPPOINT--port_number-----<<
```

Parameters

port_number

Specifies the port number of the server. Valid values are 1024 - 32767. There is no default.

Examples

TCPADMINPORT

The TCPADMINPORT option specifies the port number on which the server TCP/IP communication driver waits for requests for TCP/IP and SSL-enabled sessions other than client sessions. This includes administrative sessions, server-to-server sessions, storage agent sessions, library client sessions, managed server sessions, and event server sessions.

Using different port numbers for the options TCPPORT and TCPADMINPORT enables you to create one set of firewall rules for client sessions and another set for the previously listed session types. By using the SESSIONINITIATION parameter of REGISTER NODE and UPDATE NODE commands, you can close the port specified by TCPPORT at the firewall, and specify nodes whose scheduled sessions will be started from the server. If the two port numbers are different, separate threads are used to service client sessions and the session types. If you allow the two options to use the same port number (by default or by explicitly setting them to the same port number), a single server thread is used to service all session requests.

Client sessions attempting to use the port specified by TCPADMINPORT are terminated (if TCPPORT and TCPADMINPORT specify different ports). Administrative sessions are allowed on either port, (unless the ADMINONCLIENTPORT option is set to NO) but by default administrative sessions use the port that is specified by TCPADMINPORT.

SSL-enabled sessions that use the TCPADMINPORT option have the same limitations as the SSLTCPADMINPORT option. The following types of sessions do not use the Secure Sockets Layer (SSL) protocol:

- Network Data Management Protocol (NDMP)
- Automated Cartridge System Library Software (ACSLs)
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require that different port numbers are specified for the TCPADMINPORT and TCPPORT options.

Syntax

```
>>-TCPADMINPort--port_number-----<<
```

Parameters

port_number

Specifies the port number of the server. Valid values are 1024 - 32767. The default is the value of TCPPORT.

Examples

```
tcpadminport 1502
```

AIX | Linux

TCPBUFSIZE

The TCPBUFSIZE option specifies the size of the buffer used for TCP/IP send requests. During a restore, client data moves from the IBM Spectrum Protect™ session component to a TCP communication driver. The TCPBUFSIZE option determines if the server sends the data directly from the session buffer or copies the data to the TCP buffer. A 32 KB buffer size forces the server to copy data to its communication buffer and flush the buffer when it fills.

Note: This option is not related to the TCPWINDOWSIZE option.

Syntax

```
>>-TCPBufsize--kilobytes-----<<
```

Parameters

kilobytes

Specifies the size, in kilobytes, of the buffer used for TCP/IP send requests.

AIX The value range is from 1 to 64. The default is 32.

Linux The value range is from 1 to 64. The default is 16.

Examples

```
tcpbufsize 5
```

TCPNODELAY

The TCPNODELAY option specifies whether the server disables the delay of sending successive small packets on the network.

Change the value from the default of YES only under one of these conditions:

- You are directed to change the option by your service representative.
- You fully understand the effects of the TCP Nagle algorithm on network transmissions. Setting the option to NO enables the Nagle algorithm, which delays sending small successive packets.

Syntax

```
>>-TCPNodeLay--+-Yes-+-----><
                '-No--'
```

Parameters

Yes

Specifies that the server allows successive small packets to be sent immediately over the network. Setting this option to YES might improve performance in some high-speed networks. The default is YES.

No

Specifies that the server does not allow successive small packets to be sent immediately over the network.

Examples

```
tcpnodelay no
```

TCPPORT

The TCPPORT option specifies the port number on which the server TCP/IP communication driver waits for requests for client sessions. The server TCP/IP communication driver listens on this port for both TCP/IP and SSL-enabled sessions from the client.

Using different port numbers for the options TCPPORT and TCPADMINPORT enables you to create one set of firewall rules for client sessions and another set for other session types (administrative sessions, server-to-server sessions, storage agent sessions, library client sessions, managed server sessions, and event server sessions). If the two port numbers are different, separate threads are used to service client sessions and the other session types. If you allow the two options to use the same port number (by default or by explicitly setting them to the same port number), a single server thread is used to service all session requests.

SSL-enabled client sessions that use the TCPPORT option have the same limitations as the SSLTCPPORT option. The following types of sessions do not use SSL:

- Network Data Management Protocol (NDMP)
- Automated Cartridge System Library Software (ACSLs)
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require that different port numbers are specified for the TCPADMINPORT and TCPPORT options.

If you specify the same port number for both the SSLTCPPORT and TCPPORT options, only SSL connections are accepted and TCP/IP connections are disabled for the port.

Windows You can change this option with the SETOPT command. When you change a port, the IBM Spectrum Protect™ server starts listening on the new port immediately. All current connections remain in use until closed.

Syntax

```
>>-TCPport--port_number-----<<
```

Parameters

port_number
Specifies the port number of the server. Valid values are 1024 - 32767. The default value is 1500.

tcpport 1500

TCPWINDOWSIZE

The TCPWINDOWSIZE option specifies, in kilobytes, the amount of receive data that can be buffered at one time on a TCP/IP connection. The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window lets the sender continue sending data, and may improve communication performance, especially on fast networks with high latency.

Note:

- To improve backup performance, increase the TCPWINDOWSIZE on the server. To improve restore performance, increase the TCPWINDOWSIZE on the client.
- The TCP window acts as a buffer on the network.
- A window size larger than the buffer space on the network adapter might degrade throughput due to resending packets that were lost on the adapter.
- **AIX** | **Linux** The TCPWINDOWSIZE option is not related to the TCPBUFFSIZE option nor to the send and receive buffers allocated in client or server memory.

Syntax

```
>>-TCPWindowsize--kilobytes-----<<
```

Parameters

kilobytes
Specifies the size you want to use, in kilobytes, for the TCP/IP sliding window for your client node. You can specify a value from 0 to 2048. The default is 63. If you specify 0, the server uses the default window size set by the operating system. Values from 1 to 2048 indicate that the window size is in the range of 1 KB to 2 MB.

Examples

```
tcpwindowsize 63
```

TECBEGINEVENTLOGGING

The ECBEGINEVENTLOGGING option specifies whether event logging for the Tivoli® receiver should begin when the server starts up. If the TECHOST option is specified, ECBEGINEVENTLOGGING defaults to YES.

Syntax

```
>>-TECBegineventlogging---+Yes+-----<<  
' -No-- '
```

Parameters

- Yes
Specifies that event logging begins when the server starts up and if a TECHOST option is specified.
- No
Specifies that event logging should not begin when the server starts up. To later begin event logging to the TIVOLI receiver (if the TECHOST option has been specified), you must issue the BEGIN EVENTLOGGING command.

Examples

```
tecbegineventlogging yes
```

TECHOST

The TECHOST option specifies the host name or IP address for the Tivoli® event server.

Syntax

```
>>-TECHost--host_name-----<<
```

Parameters

- host_name
Specifies the host name or IP address for the Tivoli event server.

Examples

```
techost 9.114.22.345
```

TECPORT

The TECPORT option specifies the TCP/IP port address on which the Tivoli® event server is listening. This option is only required if the Tivoli event server is on a system that does not have a Port Mapper service running.

Syntax

```
>>-TECPort--port_number-----<<
```

Parameters

- port_number
Specifies the Tivoli event server port address. The value must be between 0 and 32767. **AIX** **Linux** This option is not required.

Examples

```
tecport 1555
```

TECUTF8EVENT

The TECUTF8EVENT option allows the IBM Spectrum Protect™ administrator to send information to the Tivoli Enterprise Console® (TEC) server in UTF-8 data format. The default is No. You can display whether or not this option is enabled by issuing the QUERY OPTION command.

Syntax

```
>>-TECUTF8event---+-Yes-+-----><
      '-No--'
```

Parameters

Yes

Specifies that the IBM Spectrum Protect server will encode the TEC event into UTF-8 before issuing the event to the TEC server.

No

Specifies that IBM Spectrum Protect server will not encode the TEC event into UTF-8 and it will be issued to the TEC server in ASCII format.

Examples

```
tecutf8event yes
```

THROUGHPUTDATATHRESHOLD

The THROUGHPUTDATATHRESHOLD option specifies a throughput threshold that a client session must reach to prevent being canceled after the time threshold is reached.

This option is used in conjunction with the THROUGHPUTTIMETHRESHOLD server option, which sets the value for the time threshold plus the media wait time. The time threshold starts when the client begins sending data to the server for storage (as opposed to setup or session housekeeping data).

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-THROUGHPUTDatathreshold-- kilobytes_per_second-----><
```

Parameters

kilobytes_per_second

Specifies the throughput that client sessions must achieve to prevent cancellation after THROUGHPUTTIMETHRESHOLD minutes have elapsed. This threshold does not include time spent waiting for media mounts. A value of 0 prevents examining client sessions for insufficient throughput. Throughput is computed by adding send and receive byte counts and dividing by the length of the session. The length does not include time spent waiting for media mounts and starts at the time a client sends data to the server for storage. The default is 0. The minimum value is 0; the maximum is 99999999.

Examples

Specify that the server is to wait until 90 minutes plus the media wait time after a session has started sending data before storage examines it as a candidate for cancellation due to low throughput. If a session is not achieving 50 KB per second in transfer rates, it will be canceled.

```
throughputtimethreshold 90
Throughputdatathreshold 50
```

THROUGHPUTTIMETHRESHOLD

The THROUGHPUTTIMETHRESHOLD option specifies the time threshold for a session after which it may be canceled for low throughput.

You can update this server option without stopping and restarting the server by using the SETOPT command. See SETOPT (Set a server option for dynamic update).

Syntax

```
>>-THROUGHPUTtimethreshold--minutes-----<<
```

Parameters

minutes

Specifies the threshold for examining client sessions and canceling them if the data throughput threshold is not met (see the THROUGHPUTDATATHRESHOLD server option). This threshold does not include time spent waiting for media mounts. The time threshold starts when a client begins sending data to the server for storage (as opposed to setup or session housekeeping data). A value of 0 prevents examining client sessions for low throughput. The default is 0. The minimum value is 0; the maximum is 99999999.

Examples

Specify that the server is to wait until 90 minutes plus the media wait time after a session has started sending data before examining it as a candidate for cancellation. If a session is not achieving 50 thousand bytes per second in transfer rates, it will be canceled.

```
throughputtimethreshold 90  
Throughputdatathreshold 50
```

Windows

TIMEFORMAT

The TIMEFORMAT option specifies the format in which time is displayed by the server.

The value for the TIMEFORMAT option is overridden by the time formatting definition of the locale if the locale is successfully initialized at server startup. The locale is specified in the LANGUAGE option.

Syntax

```
>>-TIMEformat--format_number-----<<
```

Parameters

format_number

Select a number from 1 to 4 to identify the time format used by the server. The default is 1.

- | | |
|---|----------------------|
| 1 | hh:mm:ss |
| 2 | hh,mm,ss |
| 3 | hh.mm.ss |
| 4 | hh:mm:ss a.m or p.m. |
| 5 | a.m or p.m. hh:mm:ss |

Examples

```
timeformat 4
```

TXNGROUPMAX

The TXNGROUPMAX option specifies the number of objects that are transferred as a group between a client and the server between transaction commit points. The minimum value is 4 objects and the maximum value is 65000 objects. The default value is 4096 objects. The objects transferred are actual files, directories, or both. The server counts each file or directory as one object.

It is possible to affect the performance of client backup, archive, restore, and retrieve operations by using a larger value for this option:

1. If you increase the value of the TXNGROUPMAX option by a large amount, watch for possible effects on the recovery log. A larger value for the TXNGROUPMAX option can result in increased utilization of the recovery log, as well as an increased length of time for a transaction to commit. If the effects are severe enough, they can lead to problems with operation of the server.
2. Increasing the value of the TXNGROUPMAX option can improve throughput for operations storing data directly to tape, especially when storing a large number of objects. However, a larger value of the TXNGROUPMAX option can also increase the number of objects that must be resent in the case where the transaction is stopped because an input file changed during backup, or because a new storage volume was required. The larger the value of the TXNGROUPMAX option, the more data must be resent.
3. Increasing the TXNGROUPMAX value will affect the responsiveness of stopping the operation and the client may have to wait longer for the transaction to complete.

You can override the value of this option for individual client nodes. See the TXNGROUPMAX parameter in REGISTER NODE (Register a node) and UPDATE NODE (Update node attributes).

This option is related to the TXNBYTELIMIT option in the client options file. TXNBYTELIMIT controls the number of bytes, as opposed to the number of objects, that are transferred between transaction commit points. At the completion of transferring an object, the client commits the transaction if the number of bytes transferred during the transaction reaches or exceeds the value of TXNBYTELIMIT, regardless of the number of objects transferred.

Syntax

```
>>-TXNGroupmax--number_of_objects-----<<
```

Parameters

number_of_objects

Specifies a number from 4 to 65000 for the maximum number of objects per transaction. The default is 4096.

Examples

```
txngroupmax 4096
```

UNIQUETDPTECEVENTS

The UNIQUETDPTECEVENTS option generates a unique Tivoli Enterprise Console® (TEC) event class for each individual IBM Spectrum Protect™ message, including client, server, and IBM Spectrum Protect Data Protection client messages. The default is No.

Syntax

```
>>-UNIQUETDPtecevents--+Yes+-----<<  
      '-No--'
```

Parameters

Yes

Specifies that unique IBM Spectrum Protect Data Protection messages are sent to the TEC event server. Dynamically sets UNIQUETEEvents to YES.

No

Specifies that general messages are sent to the TEC event server.

Examples

```
uniquetdptecevents yes
```

UNIQUETECEVENTS

The UNIQUETECEVENTS option generates a unique Tivoli Enterprise Console® (TEC) event class for each individual IBM Spectrum Protect™ message. The default is No.

Syntax

```
>>-UNIQUETECEvents---+-Yes-+----->>  
                        '-No--'
```

Parameters

- Yes
Specifies that unique messages are sent to the TEC event server.
- No
Specifies that general messages are sent to the TEC event server.

Examples

```
uniquetecevents yes
```

USEREXIT

The USEREXIT option specifies a user-defined exit that will be given control to manage an event.

Syntax

```
>>-USEREXIT---+-Yes-+---(1) (2)----->  
                        'module_name-----DLL_name-----'  
                        '-No--'  
  
                        (3)  
>--function-----><
```

Notes:

1. *module_name* is available only on AIX, HP-UX, Linux, Solaris, and z/OS.
2. *DLL_name* is available only on Windows.
3. *function* is available only on Windows.

Parameters

- Yes
Specifies that event logging to the user exit receiver begins automatically at server startup.
- No
Specifies that event logging to the user exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the BEGIN EVENTLOGGING command.

AIX | **Linux** *module_name*

AIX | **Linux** Specifies the module name of the user exit.

AIX | **Linux** This is the name of a shared library containing the exit. The module name can be either a fully qualified path name or just the module name itself. If it is just the module name, it is loaded from the current directory.

Windows *DLL_name*

Windows Specifies the DLL name that contains the user-exit function.

Windows function

Windows Specifies the name of the user-exit function in the DLL.

Examples

Windows

```
userexit yes dllname.dll dllmodulename
```

AIX | **Linux**

```
userexit yes fevent.exit
```

VERBCHECK

The VERBCHECK option specifies that the server will do additional error checking on the structure of commands sent by the client. This option should only be enabled when the client sends incorrectly formed requests to the server, causing the server to crash. When this option is enabled, you will get a protocol error instead of a server crash.

Syntax

```
>>-VERBCHECK-----<<
```

Parameters

None

Examples

Enable additional error checking for commands sent by the client:

```
verbcheck
```

VOLUMEHISTORY

The VOLUMEHISTORY option specifies the name of files to be automatically updated whenever server sequential volume history information is changed. There is no default for this option.

You can include one or more VOLUMEHISTORY options in the server options file. When you use multiple VOLUMEHISTORY options, the server automatically updates and stores a backup copy of the volume history information in each file you specify.

Syntax

```
>>-VOLUMEHistory--file_name-----<<
```

Parameters

file_name

Specifies the name of the file where you want the server to store a backup copy of the volume history information that it collects.

Examples

```
volumehistory volhist.out
```

Server utilities

Use server utilities to perform special tasks on the server while the server is not running.

- **Windows** DSMMAXSG (Increase the block size for writing data)
Use the DSMMAXSG utility to increase the maximum transfer length for host bus adapters (HBAs). As a result, the block size that is used by the IBM Spectrum Protect™ server for writing data to and getting data from certain types of tape drives is increased.
- DSMSERV (Start the server)
Use this utility to start the IBM Spectrum Protect server.
- **AIX** | **Linux** Server startup script: rc.dsmserv
You can use the rc.dsmserv script in your system startup to automatically start a server instance under a specific user ID.
- **Linux** Server startup script: dsmserv.rc
You can use the dsmserv.rc script to stop a server instance, or to manually or automatically start a server.
- DSMSERV DISPLAY DBSPACE (Display information about database storage space)
Use this utility to display information about storage space that is defined for the database. The output of this utility is the same as the output of the QUERY DBSPACE command, but you can use this utility when the server is not running.
- DSMSERV DISPLAY LOG (Display recovery log information)
Use this utility to display information about recovery logs including the active log, the mirror for the active log, the failover directory for the archive log, and the overflow location for logs. Use this utility when the server is not running.
- DSMSERV EXTEND DBSPACE (Increase space for the database)
Use this utility to increase space for the database by adding directories for the database to use. This utility performs the same function as the EXTEND DBSPACE command, but you can use it when the server is not running.
- DSMSERV FORMAT (Format the database and log)
Use the DSMSERV FORMAT utility to initialize the server database and recovery log. No other server activity is allowed while initializing the database and recovery log.
- DSMSERV INSERTDB (Move a server database into an empty database)
Use the DSMSERV INSERTDB utility to move a server database into a new database. The database can be extracted from the original server and inserted into a new database on the new server by using a network connection between the two servers. The database can also be inserted from media that contains the extracted database.
- DSMSERV LOADFORMAT (Format a database)
Use the DSMSERV LOADFORMAT utility when upgrading from Version 5. The utility formats an empty database in preparation for inserting an extracted database into the empty database.
- DSMSERV REMOVEDB (Remove a database)
Use the DSMSERV REMOVEDB utility to remove an IBM Spectrum Protect server database.
- DSMSERV RESTORE DB (Restore the database)
Use this utility to restore a database by using a database backup.
- **Windows** DSMSERV UPDATE (Create registry entries for a server instance)
Use this utility to create registry entries for an IBM Spectrum Protect server instance if the entries were accidentally deleted.
- **AIX** | **Linux** DSMULOG (Capture IBM Spectrum Protect server messages to a user log file)
Use this command to capture IBM Spectrum Protect server console messages to a user log file. You can specify that IBM Spectrum Protect writes messages to more than one user log file.

Windows

DSMMAXSG (Increase the block size for writing data)

Use the DSMMAXSG utility to increase the maximum transfer length for host bus adapters (HBAs). As a result, the block size that is used by the IBM Spectrum Protect™ server for writing data to and getting data from certain types of tape drives is increased.

With this utility, the maximum block size that you can specify is 256 KB. Depending on your system environment, increasing the block size can improve the rate at which IBM Spectrum Protect processes data for backup and restore operations and for archive and retrieve operations. However, the utility does not affect the generation of backup sets.

You can use tape drives that are only attached to SCSI or Fibre Channel HBAs and that have the following device types:

- 3590
- 3592
- DLT
- ECARTRIDGE
- LTO

The utility runs automatically as part of the IBM Spectrum Protect server and storage agent installation. However, if you install a new HBA on your system after you install a server or storage agent, or if you install a new version of an existing HBA device driver that resets the value of the maximum transfer size, you must run the utility manually to take advantage of the larger block size.

Parameters

AIX Linux `-u user_name`

AIX Linux Specifies a user name to switch to before you start the server. To start the server from the root user ID, you must specify the `-u` parameter and follow the instructions in Starting the server from the root user ID.

AIX Linux `-i instance_dir`

AIX Linux Specifies an instance directory to use. The instance directory becomes the current working directory of the server.

Windows `-k key_name`

Windows Specifies the name of the Windows registry key from which to retrieve information about the server. The default is `Server1`.

AIX Linux `-noexpire`

AIX Linux Specifies that the server does not remove expired files from the server database. The files are not deleted from server storage when you start the server.

Windows `NOEXPIRE`

Windows Specifies that the server does not remove expired files from the server database. The files are not deleted from server storage when you start the server.

`-o options_file`

Specifies an options file to use.

AIX Linux `-quiet`

AIX Linux Specifies that messages to the console are suppressed.

AIX Linux Windows `MAINTenance`

AIX Linux Windows Specifies that the server is started in maintenance mode, and that administrative command schedules, client schedules, client sessions, storage-space reclamation, inventory expiration, and storage-pool migration are disabled.

Tip: Maintenance mode is the preferred method for running the server during maintenance or reconfiguration tasks. When you run the server in maintenance mode, operations that might disrupt maintenance or reconfiguration tasks are disabled automatically.

`RUNFILE file_name`

Specifies the name of a text file to be run on the server. The file contains a list of server commands.

Attention: Whenever the `RUNFILE` parameter is used, the server halts when processing is complete. You must restart the server by using the `DSMSERV` utility.

Example: Start the server

Start the server for normal operation. Issue the following command on one line:

AIX

```
LDR_CNTRL=TEXTPSIZE=64K@DATAPSIZE=64K@STACKPSIZE=64K@SHMPSIZE=64K
usr/bin/dsmserv
```

AIX

Ensure that you include a space after `SHMPSIZE=64K`. By starting the server with this command, you enable 64 KB memory pages for the server. This setting helps you optimize server performance.

Linux

```
/opt/tivoli/tsm/server/bin/dsmserv
```

Windows

```
C:\Program Files\Tivoli\TSM\bin\dsmserv -k server2
```

Windows

Example: Start an additional server

Start an additional server by using the registry key named `SERVER2`.

```
dsmserv -k server2
```

AIX Linux Windows

Example: Load the sample script

Load the sample script file that is provided with the server.

```
dsmserv runfile scripts.smp
```

AIX

Linux

Windows

Example: Start the server in maintenance mode

Before you begin maintenance or reconfiguration tasks, start the server in maintenance mode.

```
dsmserv maintenance
```

Related tasks:

Starting the server in maintenance mode

AIX

Server startup script: rc.dsmserv

You can use the rc.dsmserv script in your system startup to automatically start a server instance under a specific user ID.

Syntax

```
>>-rc.dsmserv--+- -u--user_name+---+-----+-----><
      '- -U--user_name-' '- -i--instance_dir-'
```

Parameters

-u user_name

Specifies the instance user ID for which the environment is set up. The server will run under this user ID.

-U user_name

Specifies the instance user ID for which the environment is set up. The server will run under the user ID of the invoker of the command.

-i instance_dir

Specifies an instance directory, which becomes the working directory of the server.

Related tasks:

[AIX: Automatically starting servers](#)

Linux

Server startup script: dsmserv.rc

You can use the dsmserv.rc script to stop a server instance, or to manually or automatically start a server.

Prerequisites

Before you issue the DSMSERV.RC command, complete the following steps:

1. Ensure that the server instance runs under a non-root user ID with the same name as the instance owner.
2. Copy the dsmserv.rc script to the /etc/rc.d/init.d directory. The dsmserv.rc script is in the server installation directory, for example, /opt/tivoli/tsm/server/bin.
3. Rename the script so that it matches the name of the server instance owner, for example, tsminst1.
4. If the server instance directory is not home_directory/tsminst1, locate the following line in the script copy:

```
instance_dir="${instance_home}/tsminst1"
```

Change the line so that it points to your server instance directory, for example:

```
instance_dir="/tsminst1"
```

5. In the script copy, locate the following line:

```
# pidfile: /var/run/dsmserv_instancename.pid
```

Change the instance name value to the name of the server instance owner. For example, if the server instance owner is tsminst1, update the line as shown:

```
# pidfile: /var/run/dsmserv_tsminst1.pid
```

6. Use tools such as the CHKCONFIG utility to configure the run level in which the server automatically starts. Specify a value that corresponds to a multiuser mode, with networking turned on. Typically, the run level to use is 3 or 5, depending on the operating system and its configuration. For details about run levels, see the documentation for your operating system.

Syntax

```
>>-dsm serv.rc----->>
      +-start---+
      +-stop----+
      +-status---+
      '-restart-'
```

Parameters

start
Starts the server.

stop
Stops the server.

status
Shows the status of the server. If the status is *started*, the process ID of the server process is also shown.

restart
Stops the server and starts it again.

Related tasks:

[Linux: Automatically starting servers on Linux systems](#)

DSMSERV DISPLAY DBSPACE (Display information about database storage space)

Use this utility to display information about storage space that is defined for the database. The output of this utility is the same as the output of the QUERY DBSPACE command, but you can use this utility when the server is not running.

Syntax

```
>>-DSMSERV +----->
          | (1) |
          '----- -u--user_name-'

          (2) .- -k--Server1--.
>--+----->
      | (1) |      '- -k--key_name-'
      '----- -i--instance_dir-'

>--+----->
      '- -o--options_file-' '- -noexpire-' '- -quiet-'

>--DISPlay DBSPace----->>
```

Notes:

1. This parameter applies to AIX® and Linux servers only.
2. This parameter applies only to Windows servers.

Parameters

AIX | **Linux** -u user_name

- AIX** | **Linux** Specifies a user name to switch to before initializing the server.
- AIX** | **Linux** `-i instance_dir`
- AIX** | **Linux** Specifies an instance directory to use. This becomes the current working directory of the server.
- Windows** `-k key_name`
- Windows** Specifies the name of a Windows registry key that is used to store information about this server. Use this parameter only when there is more than one server on the same system. The default value is SERVER1.
- `-o options_file`
Specifies an options file to use.
- `-noexpire`
Specifies that expiration processing is suppressed when starting.
- `-quiet`
Specifies that messages to the console are suppressed.

Example: Display database space information

Display information about database storage space. See Field descriptions for details about the information shown in the output. Issue the command.

```
dsmserv display dbspace
```

Location	Total Space (MB)	Used Space (MB)	Free Space (MB)
/tsmdb001	46,080.00	20,993.12	25,086.88
/tsmdb002	46,080.00	20,992.15	25,087.85

Location	Total Space (MB)	Used Space (MB)	Free Space (MB)
d:\tsm\db001	46,080.00	20,993.12	25,086.88
d:\tsm\db002	46,080.00	20,993.15	25,087.85

Field descriptions

Location

The directory or path that is used for storing the database

Total Space (MB)

The total number of megabytes in the location

Used Space (MB)

The number of megabytes in use in the location

Free Space (MB)

The space remaining in the file system where the path is located

The space remaining on the drive where the directory is located

DSMSERV DISPLAY LOG (Display recovery log information)

Use this utility to display information about recovery logs including the active log, the mirror for the active log, the failover directory for the archive log, and the overflow location for logs. Use this utility when the server is not running.

Syntax

```
>>-DSMSERV +-----+----->
          | (1)          |
          |----- -u--user_name-|
          |
          | (2) .- -k--Server1--.
>+-----+----->
          | (1)          |          |----- -k--key_name-|
          |----- -i--instance_dir-|
>+-----+-----+-----+----->
          | -o--options_file-| | -noexpire-| | -quiet-|
```


>--DISPLAY LOG-----<<

Notes:

1. This parameter applies to AIX® and Linux servers only.
2. This parameter applies only to Windows servers.

Parameters

AIX	Linux	-u user_name	
			Specifies a user name to switch to before initializing the server.
AIX	Linux	-i instance_dir	
			Specifies an instance directory to use. This becomes the current working directory of the server.
Windows		-k key_name	
			Specifies the name of the Windows registry key from which to retrieve information about the server. Use this parameter only when there is more than one server on the same system. The default is SERVER1.
		-o options_file	
			Specifies an options file to use.
		-noexpire	
			Specifies that expiration processing is suppressed when starting.
		-quiet	
			Specifies that messages to the console are suppressed.

Examples: Display recovery log information

Display information about the recovery logs. See Field descriptions for details about the information shown in the output.

```
dsmserve display log
```

AIX	Linux
	Total Space (MB): 38,912
	Used Space (MB): 401.34
	Free Space (MB): 38,358.65
	Active Log Directory: /activelog
	Archive Log Directory: /archivelog
	Mirror Log Directory: /mirrorlog
	Archive Failover Log Directory: /archfailoverlog
Windows	
	Total Space (MB): 38,912
	Used Space (MB): 401.34
	Free Space (MB): 38,358.65
	Active Log Directory: h:\tsm\activelog
	Archive Log Directory: k:\tsm\archivelog
	Mirror Log Directory: i:\tsm\mirrorlog
	Archive Failover Log Directory: j:\tsm\archfailoverlog

Field descriptions

Total Space

Specifies the maximum size of the active log.

Used Space

Specifies the total amount of active log space currently used in the database, in megabytes.

Free Space

Specifies the amount of active log space in the database that is not being used by uncommitted transactions, in megabytes.

Active Log Directory

Specifies the location where active log files are stored. When you change the active log directory, the server moves all archived logs to the archive log directory and all active logs to a new active log directory.

Mirror Log Directory

Specifies the location where the mirror for the active log is maintained.

Archive Failover Log Directory

Specifies the location in which the server saves archive logs if the logs cannot be archived to the archive log destination.


```
>-----<
'-MIRRORlogdirectory-----directory-'
```

Notes:

1. This parameter applies to AIX® and Linux servers only.
2. This parameter applies only to Windows servers.

Parameters

AIX Linux -u user_name

AIX Linux Specifies a user name to switch to before initializing the server. This parameter is optional.

AIX Linux -i instance_dir

AIX Linux Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

Windows -k key_name

Windows Specifies the name of a Windows registry key that is used to store information about this server. Use this parameter only to install extra servers on the same system. After you install a server by using this parameter, you must always start it with the value of this parameter. This parameter is optional. The default is SERVER1.
Restriction: Additional instances of the IBM Spectrum Protect™ server that are running on the same system will compete for resources and impact overall performance of each IBM Spectrum Protect server.

-o options_file

Specifies an options file to use. This parameter is optional.

-noexpire

Specifies that expiration processing is suppressed when starting. This parameter is optional.

-quiet

Specifies that messages to the console are suppressed. This parameter is optional.

DBDir

Specifies the relative path names of one or more directories that are used to store database objects. Directory names must be separated by commas but without spaces. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

DBFile

Specifies the name of a file that contains the relative path names of one or more directories that are used to store database objects. Each directory name must be on a separate line in the file. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

ACTIVELOGSize

Specifies the size of the active log file in megabytes. This parameter is optional. The minimum value is 2048 MB (2 GB); the maximum is 524,288 MB (512 GB). If an odd number is specified, the value is rounded up to the next even number. The default is 16384 MB.

The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

ACTIVELOGDirectory (Required)

Specifies the directory in which the server writes and stores active log files. There is only one active log location. The name must be a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. The maximum number of characters is 175.

ARCHLogdirectory (Required)

Specifies the directory for the archive log files. The name must be a fully qualified directory name. The maximum number of characters is 175.

ARCHFailoverlogdirectory

Specifies the directory to be used as an alternative storage location if the ARCHLOGDIRECTORY directory is full. This parameter is optional. The maximum number of characters is 175.

MIRRORlogdirectory

Specifies the directory in which the server mirrors the active log (those files in the ACTIVELOGDIRECTORY directory). This parameter is optional. The directory must be a fully qualified directory name. The maximum number of characters is 175.

Example: Format a database

AIX Linux

```
dsmserv format dbdir=/tsmdb001 activelogsiz=8192
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

Windows

```
dsmserv -k server2 format dbdir=d:\tsm\db001 activelogsiz=8192
activelogdirectory=e:\tsm\activelog archlogdirectory=f:\tsm\archlog
archfailoverlogdirectory=g:\tsm\archfaillog mirrorlogdirectory=h:\tsm\mirrorlog
```

DSMSERV INSERTDB (Move a server database into an empty database)

Use the DSMSERV INSERTDB utility to move a server database into a new database. The database can be extracted from the original server and inserted into a new database on the new server by using a network connection between the two servers. The database can also be inserted from media that contains the extracted database.

Before you use the DSMSERV INSERTDB utility, complete the planning and preparation tasks, such as backing up the database and saving configuration information. Ensure that you meet all requirements before you move the server database.

Requirements for insertion by using media

Before you run the utility to insert the server database into an empty database, ensure that your system meets the following requirements.

- The manifest file from the DSMUPGRD EXTRACTDB operation must be available.
- If the manifest file does not contain device configuration information, or if you are specifying the CONFIGINFO=DEVCONFIG parameter, both of the following statements must be true:
 - The server options file must contain an entry for the device configuration file.
 - The device configuration file must have information about the device class that is specified in the manifest file.
- The media that contains the extracted database must be available to the V8 server. Also, the permissions must be set to grant access to the media for the user ID that owns the V8 server instance.

Syntax

```
>>-DSMSERV -+-----+----->
          | (1) |
          '----- -u--user_name-'
(2) .- -k--Server1--.
>-+-----+-----+----->
  | (1) | | | '- -k--key_name-'
  '----- -i--instance_dir-'
>-+-----+-----+-----+----->
  '- -o--options_file-' '- -noexpire-' '- -quiet-'
>>-INSERTDB--+| A: Insert from media |----->
          '-| B: Insert over a network |-'
.-PREview----No-----.
>-+-----+-----+----->>
  '-PREview----+Yes+-'
          '-No--'
A: Insert from media
|----->
```

```
'-DEVclass-----device_class_name-'
.-CONFiginfo-----MANifest-----
>-----+-----MANifest-----file_name-----|
'-CONFiginfo-----+MANifest--+-'
'-DEVconfig-'
```

B: Insert over a network

```
.-SESSWait-----60-----
|-----+-----|
'-SESSWait-----minutes-'
```

Notes:

1. This parameter applies to AIX® and Linux servers only.
2. This parameter applies only to Windows servers.

Parameters

AIX | **Linux** **-u user_name**
AIX | **Linux** Specifies a user name to switch to before initializing the server. This parameter is optional.

AIX | **Linux** **-i instance_dir**
AIX | **Linux** Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

Windows **-k key_name**
Windows Specifies the name of the Windows registry key from which to retrieve information about the server. This parameter is optional. The default is SERVER1.

-o options_file
 Specifies an options file to use. This parameter is optional.

-noexpire
 Specifies that expiration processing is suppressed when starting. This parameter is optional.

-quiet
 Specifies that messages to the console are suppressed. This parameter is optional.

DEVclass
 Specifies a sequential-access device class. You can specify any device class except for the DISK device class. The definition for the device class must exist in either the manifest file or the device configuration file.
 This parameter is optional and is used only when the database that you want to insert into the empty V8 database was extracted to media. If the database is on media and you do not specify a device class, the device class that is identified in the manifest file is used.
 Restriction: You cannot use a device class with a device type of NAS or CENTERA.

MANifest
 Specifies the location of the manifest file. Use a fully qualified file name, or place in a local directory. For example:
 ./manifest.txt
 This parameter is required when the database that you want to insert into the empty V8 database was extracted to media.

CONFiginfo
 Specifies the source of the device configuration information that is used by the DSMSERV INSERTDB operation. The default value for this parameter is MANIFEST. Possible values are as follows:

MANifest
 Specifies that device configuration information is read from the manifest file. If the manifest file does not have device configuration information, the device configuration file is used instead.

DEVConfig
 Specifies that device configuration information is read from the device configuration file.

SESSWait
 Specifies the number of minutes that the V8 server waits to be contacted by the original server. The default value is 60 minutes.
 Use this parameter only if the data that is inserted into the empty V8 database is transmitted from the source server with a network connection.

PREview

-quiet

Specifies that messages to the console are suppressed. This parameter is optional.

DBDir

Specifies the relative path names of one or more directories that are used to store database objects. Directory names must be separated by commas but without spaces. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

DBFile

Specifies the name of a file that contains the relative path names of one or more directories that are used to store database objects. Each directory name must be on a separate line in the file. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

ACTIVELOGSize

Specifies the size of the active log file in megabytes. This parameter is optional. The minimum value is 2048 MB (2 GB); the maximum is 524,288 MB (512 GB). If an odd number is specified, the value is rounded up to the next even number. The default is 16384 MB.

The size of an active log file is based on the value of the ACTIVELOGSIZE option. Guidelines for space requirements are in the following table:

Table 1. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

ACTIVELOGDirectory (Required)

Specifies the directory in which the server writes and stores active log files. There is only one active log location. The name must be a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. The maximum number of characters is 175.

ARCHLogdirectory (Required)

Specifies the directory for the archive log files. The name must be a fully qualified directory name. The maximum number of characters is 175.

ARCHFailoverlogdirectory

Specifies the directory to be used as an alternative storage location if the ARCHLOGDIRECTORY directory is full. This parameter is optional. The maximum number of characters is 175.

MIRRORlogdirectory

Specifies the directory in which the server mirrors the active log (those files in the ACTIVELOGDIRECTORY directory). This parameter is optional. The directory must be a fully qualified directory name. The maximum number of characters is 175.

Example: Format a database

AIX Linux

```
dmserv loadformat dbdir=/tsmdb001 activelogsiz=8192
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

Windows

```
dmserv -k server2 loadformat dbdir=d:\tms\db001 activelogsiz=8192
activelogdirectory=e:\tms\activelog archlogdirectory=f:\tms\archlog
archfailoverlogdirectory=g:\tms\archfaillog mirrorlogdirectory=h:\tms\mirrorlog
```

DSMSERV REMOVEDB (Remove a database)

Use the DSMSERV REMOVEDB utility to remove an IBM Spectrum Protect™ server database.

When you run this utility, you delete the server database, active log files, and active log mirror files. However, the archive log files and archive log failover log files are deleted only after you start a point-in-time database restore.

You must halt the IBM Spectrum Protect server before you issue this command.

DSMSERV RESTORE DB (Restore the database)

Use this utility to restore a database by using a database backup.

Restriction: You cannot restore a server database if the release level of the server database backup is different from the release level of the server that is being restored. For example, an error occurs when you restore a Version 7.1.3 database and you are using a Version 8.1 IBM Spectrum Protect™ server.

The restore operation uses database backups created with the BACKUP DB command.

Important: After a point-in-time restore operation, issue the AUDIT VOLUME command to audit all DISK volumes and resolve any inconsistencies between database information and storage pool volumes. Before restoring the database, examine the volume history file to find out about any sequential access storage pool volumes that were deleted or reused since the point in time to which the database was restored.

- DSMSERV RESTORE DB (Restore a database to its most current state)
Use the DSMSERV RESTORE DB utility to restore a database to its most current state under certain conditions.
- DSMSERV RESTORE DB (Restore a database to a point-in-time)
Use this command to restore a database to a point in time. A volume history file and a device configuration file must be available.

DSMSERV RESTORE DB (Restore a database to its most current state)

Use the DSMSERV RESTORE DB utility to restore a database to its most current state under certain conditions.

The following conditions must be met:

- An intact volume history file is available.
- The recovery logs are available.
- A device configuration file with the applicable device information is available.

Restriction: You cannot restore a server database if the release level of the server database backup is different from the release level of the server that is being restored. For example, an error occurs when you restore a Version 7.1.3 database and you are using a Version 8.1 IBM Spectrum Protect™ server.

IBM Spectrum Protect requests volume mounts to load the most recent backup series and then uses the recovery logs to update the database to its most current state.

Snapshot database backups cannot be used to restore a database to its most current state.

Syntax

```
>>>-DSMSERV +-----+----->
          | (1) |
          |----- -u--user_name-'
          (2) .- -k--Server1--.
>--+-----+-----+----->
  | (1) | | | '- -k--key_name-'
  |----- -i--instance_dir-'
>--+-----+-----+-----+---RESTORE DB----->
  '- -o--options_file-' | (1) |
                       |----- -quiet-'
>--+-----+-----+----->
  '-RECOVerydir----directory-'
>--+-----+-----+----->
  '-ACTIVELOGDir----directory-'
          .-PREview----No-----.
>--+-----+-----+-----+----->
  '-ON-----target_directory_file-' '-PREview-----+Yes+-'
                                     '-No--'
```

```

.-RESTOREKeys-----No-----
>-----+-----+-----+-----+----->
'-RESTOREKeys-----+No---+'
      +-YES--+
      '-ONLY-'

>-----+-----+-----+-----+----->>
'-PASSWORD---password_name-'

```

Notes:

1. This parameter applies to AIX® and Linux servers only.
2. This parameter applies only to Windows servers.

Parameters

AIX | **Linux** **-u** user_name
 Specifies a user name to switch to before initializing the server.

AIX | **Linux** **-i** instance_dir
AIX | **Linux** Specifies an instance directory to use. This instance directory becomes the current working directory of the server.

Windows **-k** key_name
Windows Specifies the name of the Windows registry key from which to retrieve information about the server. The default is SERVER1.

-o options_file
 Specifies an options file to use.

AIX | **Linux** **-quiet**
AIX | **Linux** Specifies that messages to the console are suppressed.

RECOVerydir
 Specifies a directory in which to store recovery log information from the database backup media. This directory must have enough space to hold this transaction recovery information and must be an empty directory. If this parameter is not specified, the default is to the directory specified by one of the following parameters in the DSMSEV FORMAT or DSMSEV LOADFORMAT utility:

- ARCHFAILOVERLOGDIRECTORY, if specified
- ARCHLOGDIRECTORY, if ARCHFAILOVERLOGDIRECTORY is not specified

ACTIVELOGDir
 Specifies a directory in which to store the log files that are used to track the active database operations. This directory must be specified only if the intent is to switch to an active log directory different from the one that had already been configured.

On
 Specifies a file that lists the directories to which the database is restored. Specify each directory on a separate line in the file. For example, the ON parameter specifies the restorelist.txt file, which contains the following list: **AIX** | **Linux**

```

/tsmdb001
/tsmdb002
/tsmdb003

```

Windows

```

e:\tsm\db001
f:\tsm\db002
g:\tsm\db003

```

If this parameter is not specified, the original directories that were recorded in the database backup are used.
 Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

PReview
 Specifies that the volume history files be examined and that the database backup volumes from the volume history file be evaluated.

1. Which set of database backup volumes best meets the most current criteria that are specified for restore processing? The volume history information provides details about the backup series ID, the operation ID (full, incremental 1, incremental 2, and so on), the date of the database backup, and the device class. This information and the parameters that are specified in the DSMSEV RESTORE DB command determine what to use to perform the

restore. The volume history file is examined to find the most recent database backup and then to restore the data by using that backup.

2. Is self-describing data available for the selected set of database backup volumes? Cross-check the volume history information for this backup series. The reconciliation reports what the self-describing data contains compared to what was learned from the volume history entries. The cross-check involves mounting one or more of the volumes that are indicated by the volume history. Then, using the self-describing data that was included in the database backup volumes, that information is reconciled against what is in the volume history for the database backup. If the information from the volume history file is inconsistent with the self-describing data, then messages are issued to identify the problem. For example, not all values are specified and available, and no self-describing data is found.

If the volume history information is consistent with self-describing data from the database backup, a message is issued indicating that the database backup can be used for restore processing.

If the volume history information is inconsistent with the self-describing data from the database backup or if the self-describing data for the backup cannot be found, error messages are issued indicating what was checked and what was missing.

If the PREVIEW parameter is not specified or if it is set to NO, and if the volume history and self-describing data from the database backup are consistent, then the restore proceeds.

If the PREVIEW parameter is not specified or if it is set to NO, and the reconciliation and validation fail, the database restore is not performed. Make extra volumes available and referred to from the volume history file, or remove the incomplete backup series or operation so that the IBM Spectrum Protect server selects a different preferred series or operation and continues processing.

If the PREVIEW parameter is set to YES, the process performs only the evaluation of the volume history file and the reconciliation and validation against the selected database backup.

AIX	Linux	Windows	RESTOREKeys
-----	-------	---------	-------------

AIX	Linux	Windows	Specifies whether to restore the server master encryption key that is used to encrypt storage pool data when the database is restored. This parameter is optional and only applies if you are using encrypted container storage pools in a cloud environment. If the server master key is protected when the database is restored, the default is YES. If the server master key is not protected when the database is restored, the default is NO. You can specify one of the following values:
-----	-------	---------	---

No

Specifies that the server master key is not restored when the database is restored.

Yes

Specifies that the server master key is restored when the database is restored. You must specify a password with this parameter.

Only

Specifies that only the server master key is restored. The database is not restored.

AIX	Linux	Windows	PASSword
-----	-------	---------	----------

AIX	Linux	Windows	Specifies the password that is used to protect the database backup. This parameter only applies if you are using encrypted container storage pools in a cloud environment. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database. You must use a password if you specify the RESTOREKEYS=YES or RESTOREKEYS=ONLY parameter.
-----	-------	---------	---

Example: Restore the database to its most current state

Restore the database to its most current state by using the already configured active log directory.

```
dmserv restore db
```

Example: Restore the server master key without restoring the database

Restore the server master key without restoring the database by issuing the following command:

```
dmserv restore db restorekeys=only
```

DSMSERV RESTORE DB (Restore a database to a point-in-time)

Use this command to restore a database to a point in time. A volume history file and a device configuration file must be available.

Windows Specifies the name of the Windows registry key from which to retrieve information about the server. The default is SERVER1.

-o options_file

Specifies an options file to use.

AIX | **Linux** -quiet

AIX | **Linux** Specifies that messages to the console are suppressed.

TODate (Required)

Specifies the date to which to restore the database. The following values are possible:

MM/DD/YYYY

Specifies that you want to restore a database by using the last backup series that was created before this specified date.

TODAY

Specifies that you want to restore a database by using the most recent backup series that was created before today.

TODAY-numdays or -numdays

Specifies that you want to restore a database by using the most recent backup series that was created the specified number of days before the current date.

TOTime

Specifies the time of day to which to restore the database. This parameter is optional. The default is the end of the day (23:59:59). Possible values are:

HH:MM:SS

Specifies that you want to restore the database by using the last backup series that is created on or before the specified time on the date that is specified on the TODATE parameter.

NOW

Specifies that you want to restore the database by using a backup series that is created on or before the current time on the date that is specified on the TODATE parameter.

For example, if you issue the DSMSEV RESTORE DB utility at 9:00 with TOTIME=NOW, the database is restored by using the last backup series that is created on or before 9:00 on the date that is specified on the TODATE parameter.

NOW-numhours:numminutes or -numhours:numminutes

Specifies that you want to restore the database by using a backup series that is created on or before the current time minus a specified number of hours and, optionally, minutes on the date that is specified on the TODATE parameter.

For example, if you issue the DSMSEV RESTORE DB utility at 9:00 with TOTIME=NOW-3:30 or TOTIME+-3:30, the database is restored by using the last backup series that is created on or before 5:30 on the date that is specified on the TODATE parameter.

Source

Specifies whether the database is restored by using either database full and incremental backup volumes or snapshot database volumes. This parameter is optional. The default value is DBBackup. The following values are possible:

DBBackup

Specifies that the database is restored as follows:

1. Reads the volume history file to locate the database full and incremental backup volumes that are needed.
2. Requests mounts and loads the data from the database full and incremental backup volumes as required to restore the database volume to the specified time.

DBSnapshot

Specifies that the database is restored as follows:

1. Reads the volume history file to locate the snapshot database volumes that are needed,
2. Requests mounts and loads data from snapshot database volumes as required to restore the volume to the specified time.

RECOVdir

Specifies a directory in which to store recovery log information from the database backup media. This log information is used to establish transaction consistency of the server database as part of the recovery processing. This directory must have enough space to hold this transaction recovery information and must be an empty directory. If this parameter is not specified, the default is the directory that is specified by one of the following parameters in the DSMSEV FORMAT or DSMSEV LOADFORMAT utility:

- ARCHFAILOVERLOGDIRECTORY, if specified
- ARCHLOGDIRECTORY, if ARCHFAILOVERLOGDIRECTORY is not specified

ACTIVELOGDir

Specifies a directory in which to store the log files that are used to track the active database operations. Specify this directory only if the intent is to switch to an active log directory that is different from the one that was already configured.

On

Specifies a file that lists the directories to which the database is restored. Specify each directory on a separate line in the file. For example, the ON parameter specifies the restorelist.txt file, which contains the following list:

```
/tsmdb001
/tsmdb002
/tsmdb003
```

Windows

```
e:\tsm\db001
f:\tsm\db002
g:\tsm\db003
```

If this parameter is not specified, the original directories that were recorded in the database backup are used.

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

PReview

Specifies that the volume history files be examined and that the database backup volumes from the volume history file be evaluated.

1. Which set of database backup volumes best meets the point-in-time criteria that are specified for restore processing? The volume history information provides details about the backup series ID, the operation ID (full, incremental 1, incremental 2, and so on), the date of the database backup, and the device class. This information and the parameters that are specified in the DSMSEV RESTORE DB command determine what to use to perform the restore. The volume history file is examined to find the best database backup that meets the specified point-in-time criteria and then perform the restore by using that backup.
2. Is self-describing data available for the selected set of database backup volumes? Cross-check the volume history information for this backup series. The reconciliation reports what the self-describing data contains compared to what was learned from the volume history entries. The cross-check involves mounting one or more of the volumes that are indicated by the volume history. Then, using the self-describing data that was included in the database backup volumes, that information is reconciled against what is in the volume history for the database backup. If the information from the volume history file is inconsistent with the self-describing data, then messages are issued to identify the problem. For example, not all values are specified and available, and no self-describing data is found.

If the volume history information is consistent with self-describing data from the database backup, a message is issued indicating that the database backup can be used for restore processing.

If the volume history information is inconsistent with the self-describing data from the database backup or if the self-describing data for the backup cannot be found, error messages are issued indicating what was checked and what was missing.

If the PREVIEW parameter is not specified or if it is set to NO, and if the volume history and self-describing data from the database backup are consistent, then the restore proceeds.

If the PREVIEW parameter is not specified or if it is set to NO, and the reconciliation and validation fail, the database restore is not performed. Make extra volumes available and referred to from the volume history file, or remove the incomplete backup series or operation so that the IBM Spectrum Protect server selects a different preferred series or operation and continues processing.

If the PREVIEW parameter is set to YES, the process performs only the evaluation of the volume history file and the reconciliation and validation against the selected database backup.

AIX Linux Windows RESTOREKeys

Specifies whether to restore the server master encryption key that is used to encrypt storage pool data when the database is restored. This parameter is optional and only applies if you are using encrypted container storage pools in a cloud environment. If the server master key is protected when the database is restored, the default is YES. If the server master key is not protected when the database is restored, the default is NO. You can specify one of the following values:

- No
Specifies that the server master key is not restored when the database is restored.
- Yes
Specifies that the server master key is restored when the database is restored. You must specify a password with this parameter.
- Only
Specifies that only the server master key is restored. The database is not restored.

AIX | **Linux** | **Windows** **PASS**word

AIX | **Linux** | **Windows** Specifies the password that is used to protect the database backup. This parameter only applies if you are using encrypted container storage pools in a cloud environment. If you specify a password for database backup, you must specify the same password on the RESTORE DB command to restore the database. You must use a password if you specify the RESTOREKEYS=YES or RESTOREKEYS=ONLY parameter.

Example: Restore the database to a specific point in time

Restore the database to its state on May 12, 2011 at 2:25 PM.

```
dsmserv restore db todate=05/12/2011 totime=14:45
```

Example: Restore the server master key without restoring the database

Restore the server master key without restoring the database by issuing the following command:

```
dsmserv restore db restorekeys=only
```

Windows

DSMSERV UPDATE (Create registry entries for a server instance)

Use this utility to create registry entries for an IBM Spectrum Protect™ server instance if the entries were accidentally deleted.

Run this utility from the instance directory for the database (where files such as dsmserv.dsk are stored for the server). The utility re-creates the original registry entries for the server.

Syntax

```

      .- -k--Server1--.
>>-DSMSERV-+-----+-----UPDATE----->>
      '- -k--key_name-'

```

Parameters

- k key_name
Specifies the name of the Windows registry key in which to store information about the server. The default is Server1.

Example: Re-create registry entries for a server instance

Run the utility to re-create registry entries for the server instance, Server2.

```
"c:\Program Files\Tivoli\TSM\server\bin\dsmserv" -k server2 update
```

AIX | **Linux**

DSMULOG (Capture IBM Spectrum Protect server messages to a user log file)

Use this command to capture IBM Spectrum Protect™ server console messages to a user log file. You can specify that IBM Spectrum Protect writes messages to more than one user log file.

Important: Do not place the user logs in the /usr or /opt file systems because space constraints in the file system can prevent the server from starting.

Syntax

```
      .-|-----|  
      v      |  
>>-DSMULOG----logfilename+-----<<
```

Parameters

logfilename (Required)

Specifies the name of one or more user log files to which IBM Spectrum Protect writes server console messages. When you specify multiple file names, each file is written to for one day and then the server moves to the next file to capture log messages. When all the files in the list have been written to, the server begins writing to the first file again and any messages contained therein are overwritten.

Example: Capture server console messages to a user log file on a daily basis

Specify the user log files to which you want to log console messages.

In this example, if you invoke this utility on Friday, on Friday the server messages are captured to log1, on Saturday the messages are captured to log2, and on Sunday the messages are captured to log3. On Monday, the messages are captured to log1 and the messages from the previous Friday are overwritten.

```
/opt/tivoli/tsm/server/bin/dsmserv -u tsminst1 -i  
/tsmserv/tsminst1/tsminst1 2>&1 | dsmulog /tsmserv/tsminst1/tsminst1/log1  
/tsmserv/tsminst1/tsminst1/log2  
/tsmserv/tsminst1/tsminst1/log3 &
```

IBM Spectrum Protect server device utilities

You can use device utilities for tasks that are related to configuring storage devices for the server.

Device utilities

- **AIX** dsmsanlist (Display information about devices)
- **Linux** autoconf (Auto configure devices)
- **Windows** tsmdllst (Display information about devices)

AIX | **Linux**

dsmsanlist (Display information about devices)

Use the dsmsanlist device information utility to display information about devices that are connected to the IBM Spectrum Protect™ server.

The dsmsanlist utility is part of the IBM Spectrum Protect server and IBM Spectrum Protect storage agent package. The utility is installed along with the IBM Spectrum Protect server or the IBM Spectrum Protect storage agent. By default, the utility is located in either the server/bin directory (/opt/tivoli/tsm/server/bin) or the storage agent directory (/opt/tivoli/tsm/StorageAgent/bin).

The dsmsanlist utility uses the host bus adapter (HBA) API interface to get the device information from the storage area network (SAN). Therefore, before you run the utility, ensure that the HBA API library of the HBA vendor is also installed.

You can run the dsmsanlist utility by browsing to the relevant directory (either /opt/tivoli/tsm/server/bin or /opt/tivoli/tsm/StorageAgent/bin), and then entering dsmsanlist. No additional options are available for this utility.

The dsmsanlist utility displays the following information as output:

- HBA information
- HBA port number

- Device vendor ID
- Product ID
- Device type
- Device serial number
- Port worldwide name
- IBM Spectrum Protect device name

A log file (dsmsanlist.log), which you can use for debugging purposes, is also generated by default.

AIX | Linux

Example: Display information about all devices

Display information about all devices that are connected to the IBM Spectrum Protect server:

```
dsmsanlist

root@xlinux3:/opt/tivoli/tsm/server/bin# ./dsmsanlist

*****
*       IBM Spectrum Protect       *
*       dsmsanlist Utility Program   *
*****
Licensed Materials - Property of IBM

(C) Copyright IBM Corporation 2013. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Port #1  Vendor_ID  Product_ID      Type      Serial_Number    Port_WWN
Dev_Name
=====
IBM      ULTRIUM-TD8     Tape           C3EAC62000      500308c3eac62001
/dev/sg13;
QUANTUM  Scalar i3-i6    Changer       QUANTUMFFC1652024_LLA  500308c3eac62001
/dev/sg28;/dev/changer-sg28;
IBM      ULTRIUM-HH7     Tape           11C1A030B5      5000e111c1a030b6
/dev/sg1;
BDT      MULTISTAK       Changer       DE68101026_LL01  5000e111c1a030b6
/dev/sg2;/dev/changer-sg2;
IBM      ULTRIUM-HH6     Tape           11C1A030BF      5000e111c1a030c0
/dev/sg3;
HPE      Ultrium 8-SCSI  Tape           9C1730D495      5001438016044f42
/dev/sg8;
HP       1x8 G2 AUTOLDR Changer       4C6140X001      5001438016044f42
/dev/sg21;/dev/changer-sg21;
IBM      ULTRIUM-TD8     Tape           C3EAC62114      500308c3eac62115
/dev/sg24;
=====
```

Linux

autoconf (Auto configure devices)

Use the autoconf utility to configure devices for use with the IBM Spectrum Protect™ server.

The autoconf utility performs the following tasks:

- Loads the driver to the kernel
- Creates the necessary files for the IBM Spectrum Protect device driver
- Creates device information files for libraries and tape devices

The autoconf utility is included in the device driver package and is installed to the /opt/tivoli/tsm/devices/bin directory.

Options

-a

Adds read and write permissions to IBM Spectrum Protect device files to allow all users access to the devices. Specify this value to configure devices if the IBM Spectrum Protect server is started by a non-root user.

-g

Adds read and write permissions to the IBM Spectrum Protect device files to allow anyone in the same group as a root user to use the devices.

-t

Enables tracing for the autoconf utility.

-?

Displays information about the autoconf utility and its parameters.

Example: Configure devices by using the autoconf utility

Run autoconf utility to configure IBM Spectrum Protect devices:

```
> /opt/tivoli/tsm/devices/bin/autoconf
```

Linux

Example: For a server that is started by a non-root user ID, configure devices by using the autoconf utility

Run autoconf to configure IBM Spectrum Protect devices. Use the a option because the server is started by a user ID that is not the root user.

```
> /opt/tivoli/tsm/devices/bin/autoconf -a
```

```
Added the read and write permissions for all users to /dev/sg4.
Added the read and write permissions for all users to /dev/sg5.
Added the read and write permissions for all users to /dev/sg6.
Added the read and write permissions for all users to /dev/sg7.
Added the read and write permissions for all users to /dev/sg8.
Added the read and write permissions for all users to /dev/sg9.
Added the read and write permissions for all users to /dev/sg10.
Added the read and write permissions for all users to /dev/sg11.
Added the read and write permissions for all users to /dev/sg12.
Added the read and write permissions for all users to /dev/sg13.
Added the read and write permissions for all users to /dev/sg14.
Added the read and write permissions for all users to /dev/sg15.
Added the read and write permissions for all users to /dev/sg16.
Added the read and write permissions for all users to /dev/sg17.
Added the read and write permissions for all users to /dev/sg18.
Added the read and write permissions for all users to /dev/sg19.
Added the read and write permissions for all users to /dev/sg20.
Added the read and write permissions for all users to /dev/sg21.
Added the read and write permissions for all users to /dev/sg22.
Added the read and write permissions for all users to /dev/sg23.
Added the read and write permissions for all users to /dev/sg24.
Added the read and write permissions for all users to /dev/sg25.
Added the read and write permissions for all users to /dev/sg26.
Added the read and write permissions for all users to /dev/sg27.
Added the read and write permissions for all users to /dev/sg28.
Added the read and write permissions for all users to /dev/sg29.
```

Tape Drives:

=====

Index	Minor	Host	CHN	ID	LUN	Type	Vendor_ID	Device_Serial_Number	Product_ID	Rev.
000	004	003	000	004	000	001	IBM	1068000439	ULTRIUM-HH5	C5X1
001	007	003	000	008	001	001	HP	01UbWSD-04	Ultrium 2-SCSI	R210
002	008	003	000	008	002	001	HP	01UbWSD-05	Ultrium 2-SCSI	R210
003	010	003	000	008	004	001	HP	01UbWSD-07	Ultrium 3-SCSI	R210
004	012	003	000	008	006	001	HP	01UbWSD-01	Ultrium 3-SCSI	R210
005	013	003	000	008	007	001	HP	01UbWSD-02	Ultrium 3-SCSI	R210
006	014	003	000	008	008	001	HP	01UbWSD-08	Ultrium 3-SCSI	R210
007	015	003	000	008	009	001	HP	01UbWSD-09	Ultrium 3-SCSI	R210
008	016	003	000	008	010	001	HP	01UbWSD-0a	Ultrium 3-SCSI	R210
009	017	003	000	008	011	001	HP	01UbWSD-0b	Ultrium 3-SCSI	R210
010	018	003	000	008	012	001	HP	01UbWSD-0c	Ultrium 3-SCSI	R210
011	019	003	000	008	013	001	HP	01UbWSD-0d	Ultrium 3-SCSI	R210
012	020	003	000	005	000	001	IBM	1068000913	ULTRIUM-HH5	C5X1
013	022	003	000	009	001	001	QUANTUM	01UbWSD-0f	SDLT320	R210

014	023	003	000	009	002	001	QUANTUM	01UbWSD-0g	SDLT320	R210
015	024	003	000	009	003	001	QUANTUM	01UbWSD-0h	SDLT320	R210
016	025	003	000	009	004	001	QUANTUM	01UbWSD-0i	SDLT320	R210
017	026	003	000	006	000	001	IBM	1068001573	ULTRIUM-HH4	B5Q1
018	027	003	000	007	000	001	IBM	1068001545	ULTRIUM-HH4	B5Q1
019	028	003	000	010	000	001	HP	HU19477PAE	Ultrium 5-SCSI	I65W

Medium Changer Devices:

```
=====
```

Index	Minor	Host	CHN	ID	LUN	Type	Vendor_ID	Device_Serial_Number	Product_ID	Rev.
000	005	003	000	004	001	008	NEC	2Y11BB0023	LL-2B01	0004
001	006	003	000	008	000	008	HP	01UbWSD-03	VLS	1.00
002	009	003	000	008	003	008	HP	01UbWSD-06	ThinStor AutoLdr	T133
003	011	003	000	008	005	008	HP	01UbWSD-00	ESL E-Series	2.00
004	021	003	000	009	000	008	HP	01UbWSD-0e	MSL6000 Series	0430
005	029	003	000	010	001	008	HP	3615-0101	MSL G3 Series	1120

Windows

tsmdlst (Display information about devices)

Use the tsmdlst utility to display information about devices that are connected to the IBM Spectrum Protect™ server. You can view device names and other information about medium changer and tape devices that are controlled by the IBM Spectrum Protect device driver.

The tsmdlst utility is part of the IBM Spectrum Protect server and IBM Spectrum Protect storage agent package. By default, the utility is located in the devices installation directory (either C:\Program Files\Tivoli\TSM\server or C:\Program Files\Tivoli\TSM\StorageAgent). The tsmdlst utility uses the host bus adapter (HBA) API interface to get the device information from the storage area network (SAN). Therefore, before you run the utility, ensure that the HBA API library of the HBA vendor is also installed.

You can run the utility by browsing to the relevant directory (either C:\Program Files\Tivoli\TSM\server or C:\Program Files\Tivoli\TSM\StorageAgent), and then entering tsmdlst.exe. No additional options are available for this utility.

The tsmdlst utility displays the following information as output:

- HBA information
- HBA port number
- Device vendor ID
- Product ID
- Device type
- Device serial number
- Port worldwide name
- Operating system device name
- IBM Spectrum Protect device name

A log file (tsmdlst.log), which you can use for debugging purposes, is also generated by default.

Example: Display information about devices

Display information about medium changer and tape devices by running the tsmdlst utility:

```
tsmdlst.exe

C:\Program Files\Tivoli\TSM>tsmdlst.exe
*****
*       IBM Spectrum Protect       *
*       tsmdlst SAN Utility Program *
*****
Licensed Materials - Property of IBM

5608-E01
5608-E02

(C) Copyright International Business Machines Corp. 1990, 2011.
All rights reserved.

Port #1  Vendor_ID  Product_ID      Type      Serial_Number      Port_WWN
OS_Dev   Dev_Name
```

```

=====
=====
Tape6      IBM          ULTRIUM-HH6      Tape      11C1A030BF      5000e111c1a030c0
          mt2.0.0.5
Tape7      IBM          ULTRIUM-HH7      Tape      11C1A030B5      5000e111c1a030b6
          mt3.0.0.5
Tape8      IBM          ULTRIUM-TD8      Tape      C3EAC62114      500308c3eac62115
          mt0.0.0.5
Tape9      IBM          ULTRIUM-TD8      Tape      C3EAC62000      500308c3eac62001
          mt1.0.0.5
Changer    QUANTUM      Scalar i3-i6     Changer   QUANTUMFFC1652024_LLA  500308c3eac62001
          lb1.1.0.5
Changer    BDT          MULTISTAK        Changer   DE68101026_LL01      5000e111c1a030b6
          lb3.1.0.5
Tape      HPE          Ultrium 8-SCSI   Tape      9C1730D495      5001438016044f42
          mt5.0.0.5
Changer    HP           1x8 G2 AUTOLDR  Changer   4C6140X001      5001438016044f42
          lb5.1.0.5
=====
=====

```

Server scripts and macros for automation

You can automate common administrative tasks by creating IBM Spectrum Protect™ server scripts or administrative client macros. Server scripts are stored in the server database and can be scheduled to run with an administrative schedule command. Administrative client macros are stored as files on the administrative client. Macros cannot be distributed across servers and cannot be scheduled on the server.

- **Server scripts**
You can automate common administrative tasks with scripts that are stored in the server database. You can schedule a script for processing by using the administrative command scheduler on the server.
- **Administrative client macros**
A macro is a file that contains one or more administrative client commands. You can run a macro from the administrative client only in batch or interactive modes. Macros are stored as a file on the administrative client. Macros are not distributed across servers and cannot be scheduled on the server.

Server scripts

You can automate common administrative tasks with scripts that are stored in the server database. You can schedule a script for processing by using the administrative command scheduler on the server.

IBM Spectrum Protect™ scripts have the following capabilities and statements:

- Command parameter substitution.
- SELECT commands that you specify when the script is processed.
- Command execution control, such as PARALLEL and SERIAL processing options.
- Conditional logic flow statements. These logic flow statements include the following statements:
 - The IF clause; this clause determines how processing proceeds based on the current return code value.
 - The EXIT statement; this statement ends script processing.
 - The GOTO and LABEL statement. This statement directs logic flow to continue processing with the line that starts with the label specified.
- Comment lines.

Sample scripts are provided in the scripts.smp file. The sample scripts have an example order of execution for scheduling administrative commands.

If one of the specified commands in the script does not process successfully, the remaining commands are not processed.

- **Defining a server script**
You can define a server script line-by-line, create a file that contains the command lines, or copy an existing script.
- **Updating a script**
You can update a script to change a command line or to add a command line to a script.
- **Querying a server script to create another server script**
You can create more server scripts by querying a script and specifying the FORMAT=RAW and OUTPUTFILE parameters. You can use the resulting output as input into another script without having to create a script line by line.

- Running a server script
To process a script, use the RUN command. You can run a script that contains substitution variables by specifying them along with the RUN command.

Defining a server script

You can define a server script line-by-line, create a file that contains the command lines, or copy an existing script.

About this task

Restriction: You cannot redirect the output of a command within a server script. Instead, run the script and then specify command redirection. For example, to direct the output of script1 to the c:\temp\test.out directory, run the script and specify command redirection as in the following example:

```
run script1 > c:\temp\test.out
```

Procedure

1. Define a script with the DEFINE SCRIPT command. You can initially define the first line of the script with this command. For example:

```
define script qaixc "select node_name from nodes where platform='aix'"  
desc='Display AIX clients'
```

This example defines the script as QAIXC. When you run the script, all AIX® clients are displayed.

2. Define more lines in the script with the UPDATE SCRIPT command. For example, you want to add a QUERY SESSION command, enter:

```
update script qaixc "query session *"
```

3. Optional: You can specify a WAIT parameter with the DEFINE CLIENTACTION command. By using this parameter, you can specify that the client action must complete before the next step in the command script or macro is processed.
4. Optional: To help you determine where a problem is within a command in a script, use the ISSUE MESSAGE command.

- Running commands in parallel or serially
You have the options of running commands in a script serially, in parallel, or serially and in parallel. You can do so by using the SERIAL or PARALLEL script commands in the COMMAND_LINE parameter of DEFINE and UPDATE SCRIPT. Therefore, it is possible to run multiple commands in parallel and wait for them to complete before the next command is run.
- Continuing commands across multiple command lines
You can continue long commands across multiple command lines by specifying the continuation character (-) as the last character for a command that is continued.
- Including substitution variables in a script
You can include substitution variables in a script. Substitution variables are specified with a \$ character followed by a number that represents the position of the parameter when the script is processed.
- Including logic flow statements in a script
You can use conditional logic flow statements that are based on return codes that are issued from previous command processing. By using these logic statements, you can process your scripts according to the outcome of certain commands. You can use IF, EXIT, or GOTO (label) statements.
- Using SELECT commands in a script
An IBM Spectrum Protect™ script is one or more commands that are stored as an object in the database. You can define a script that contains one or more SELECT commands.

Running commands in parallel or serially

You have the options of running commands in a script serially, in parallel, or serially and in parallel. You can do so by using the SERIAL or PARALLEL script commands in the COMMAND_LINE parameter of DEFINE and UPDATE SCRIPT. Therefore, it is possible to run multiple commands in parallel and wait for them to complete before the next command is run.

About this task

Running commands serially in a script ensures that any preceding commands are complete before proceeding and ensures that any following commands are run serially. When a script starts, all commands are run serially until a PARALLEL command is

encountered. Multiple commands that are running in parallel and accessing common resources, such as tape drives, can run serially.

Script return codes remain the same before and after a PARALLEL command is run. When a SERIAL command is encountered, the script return code is set to the maximum return code from any previous commands that were run in parallel.

When you use server commands that support the WAIT parameter after a PARALLEL command, the behavior is as follows:

- If you specify (or use the default) WAIT=NO, a script does not wait for the completion of the command when a subsequent SERIAL command is encountered. The return code from that command reflects processing only up to the point that the command starts a background process. The final return code from the command is not available to your script.
- If you specify WAIT=YES, your script waits for the completion of the command when a subsequent SERIAL command is encountered. The return code from that command reflects processing for the entire command.

In most cases, you can use WAIT=YES on commands that are run in parallel.

Restriction: If the command starts a background process that does not have the WAIT parameter, the command is considered to be complete after the background thread is started. Therefore, the command can run only in parallel.

The following example illustrates how the PARALLEL command is used to back up, migrate, and reclaim storage pools.

```
/*run multiple commands in parallel and wait for
them to complete before proceeding*/
PARALLEL
/*back up four storage pools simultaneously*/
BACKUP STGPOOL PRIMPOOL1 COPYPOOL1 WAIT=YES
BACKUP STGPOOL PRIMPOOL2 COPYPOOL2 WAIT=YES
BACKUP STGPOOL PRIMPOOL3 COPYPOOL3 WAIT=YES
BACKUP STGPOOL PRIMPOOL4 COPYPOOL4 WAIT=YES
/*wait for all previous commands to finish*/
SERIAL
/*after the backups complete, migrate stgpools
simultaneously*/
PARALLEL
MIGRATE STGPOOL PRIMPOOL1 DURATION=90 WAIT=YES
MIGRATE STGPOOL PRIMPOOL2 DURATION=90 WAIT=YES
MIGRATE STGPOOL PRIMPOOL3 DURATION=90 WAIT=YES
MIGRATE STGPOOL PRIMPOOL4 DURATION=90 WAIT=YES
/*wait for all previous commands to finish*/
SERIAL
/*after migration completes, reclaim storage
pools simultaneously*/
PARALLEL
RECLAIM STGPOOL PRIMPOOL1 DURATION=120 WAIT=YES
RECLAIM STGPOOL PRIMPOOL2 DURATION=120 WAIT=YES
RECLAIM STGPOOL PRIMPOOL3 DURATION=120 WAIT=YES
RECLAIM STGPOOL PRIMPOOL4 DURATION=120 WAIT=YES
```

Related reference:

DEFINE SCRIPT (Define a server script)
UPDATE SCRIPT (Update a server script)

Continuing commands across multiple command lines

You can continue long commands across multiple command lines by specifying the continuation character (-) as the last character for a command that is continued.

About this task

The following example continues an SQL statement across multiple command lines:

```
/*-----*/
/* Sample continuation example */
SELECT-
* FROM-
NODE WHERE-
PLATFORM='win32'
```

When this command is processed, it runs the following command:

```
select * from nodes where platform='win32'
```

Including substitution variables in a script

You can include substitution variables in a script. Substitution variables are specified with a \$ character followed by a number that represents the position of the parameter when the script is processed.

About this task

The following example SQLSAMPLE script specifies substitution variables \$1 and \$2:

```
/*-----*/  
/* Sample substitution example */  
/* -----*/  
SELECT-  
$1 FROM-  
NODES WHERE-  
PLATFORM='$2'
```

When you run the script you must specify two values, one for \$1 and one for \$2. For example:

```
run sqlsample node_name aix
```

The command that is processed when the SQLSAMPLE script is run is the following command:

```
select node_name from nodes where platform='aix'
```

Including logic flow statements in a script

You can use conditional logic flow statements that are based on return codes that are issued from previous command processing. By using these logic statements, you can process your scripts according to the outcome of certain commands. You can use IF, EXIT, or GOTO (label) statements.

As each command is processed in a script, the return code is saved for possible evaluation before the next command is processed. The return code can be one of three severities: OK, WARNING, or ERROR. See Return codes for use in scripts for a list of valid return codes and severity levels.

- Specifying the IF clause
You can use the IF clause at the beginning of a command line to determine how processing of the script proceeds based on the current return code value. In the IF clause, you specify a return code symbolic value or severity.
- Specifying the EXIT statement
Use the EXIT statement to end script processing.
- Specifying the GOTO statement
The GOTO statement is used with a label statement. The label statement is the target of the GOTO statement. The GOTO statement directs script processing to the line that contains the label statement to resume processing from that point.

Specifying the IF clause

You can use the IF clause at the beginning of a command line to determine how processing of the script proceeds based on the current return code value. In the IF clause, you specify a return code symbolic value or severity.

About this task

The server initially sets the return code at the beginning of the script to RC_OK. The return code is updated by each processed command. If the current return code from the processed command is equal to any of the return codes or severities in the IF clause, the remainder of the line is processed. If the current return code is not equal to one of the listed values, the line is skipped.

The following script example backs up the BACKUPPOOL storage pool only if there are no sessions currently accessing the server. The backup proceeds only if a return code of RC_NOTFOUND is received:

```
/* Backup storage pools if clients are not accessing the server */  
select * from sessions  
/* There are no sessions if rc_notfound is received */  
if(rc_notfound) backup stg backuppool copypool
```


The following script example backs up the BACKUPPOOL storage pool if a return code with a severity of warning is encountered:

```
/* Backup storage pools if clients are not accessing the server */
select * from sessions
/* There are no sessions if rc_notfound is received */
if(warning) backup stg backuppool cypool
```

Specifying the EXIT statement

Use the EXIT statement to end script processing.

About this task

The following example uses the IF clause together with RC_OK to determine if clients are accessing the server. If an RC_OK return code is received, it indicates that client sessions are accessing the server. The script proceeds with the exit statement, and the backup does not start.

```
/* Back up storage pools if clients are not accessing the server */
select * from sessions
/* There are sessions if rc_ok is received */
if(rc_ok) exit
backup stg backuppool cypool
```

Specifying the GOTO statement

The GOTO statement is used with a label statement. The label statement is the target of the GOTO statement. The GOTO statement directs script processing to the line that contains the label statement to resume processing from that point.

About this task

The label statement always has a colon (:) after it and can be blank after the colon. The following example uses the GOTO statement to back up the storage pool only if there are no sessions currently accessing the server. In this example, the return code of RC_OK indicates that clients are accessing the server. The GOTO statement directs processing to the `done:` label, which contains the EXIT statement that ends the script processing:

```
/* Back up storage pools if clients are not accessing the server */
select * from sessions
/* There are sessions if rc_ok is received */
if(rc_ok) goto done
backup stg backuppool cypool
done:exit
```

Using SELECT commands in a script

An IBM Spectrum Protect™ script is one or more commands that are stored as an object in the database. You can define a script that contains one or more SELECT commands.

About this task

A script can be run from an administrative client or the server console. You can also include it in an administrative command schedule to run automatically. See Server scripts for details.

IBM Spectrum Protect is shipped with a file that contains a number of sample scripts. The file, `scripts.smp`, is in the server directory. To create and store the scripts as objects in your server's database, issue the `DSMSERV RUNFILE` command during installation:

```
> dsmserv runfile scripts.smp
```

You can also run the file as a macro from an administrative command line client:

```
macro scripts.smp
```

The sample scripts file contains commands. These commands first delete any scripts with the same names as those to be defined, then define the scripts. The majority of the samples create SELECT commands, but others do such things as back up storage

pools. You can also copy and change the sample scripts file to create your own scripts.

Here are a few examples from the sample scripts file:

```
def script q_inactive_days /* -----*/
upd script q_inactive_days /* Script Name: Q_INACTIVE */
upd script q_inactive_days /* Description: Display nodes that have not */
upd script q_inactive_days /* accessed the backup server for a */
upd script q_inactive_days /* specified number of days */
upd script q_inactive_days /* Parameter 1: days */
upd script q_inactive_days /* Example: run q_inactive_days 5 */
upd script q_inactive_days /* -----*/
upd script q_inactive_days "select node_name,lastacc_time from nodes where -"
upd script q_inactive_days " cast((current_timestamp-lastacc_time)days as -"
upd script q_inactive_days " decimal) >= $1 "

/* Display messages in the activity log of severity X or Y */

def script q_msg_sev desc='Show msgs in the activity log of severity X or Y'
upd script q_msg_sev /* -----*/
upd script q_msg_sev /* Script Name: Q_MSG_SEV */
upd script q_msg_sev /* Description: Display messages in the */
upd script q_msg_sev /* activity log that have either */
upd script q_msg_sev /* of two specified severities. */
upd script q_msg_sev /* Parameter 1: severity 1 */
upd script q_msg_sev /* Parameter 2: severity 2 */
upd script q_msg_sev /* where severity is I, W, E, S, or D */
upd script q_msg_sev /* Example: run q_msg_sev S E */
upd script q_msg_sev /* -----*/
upd script q_msg_sev "select date_time,msgno,message from actlog -"
upd script q_msg_sev " where severity=upper('$1') or severity=upper('$2')"
```

Updating a script

You can update a script to change a command line or to add a command line to a script.

- **Appending a new command**
To append a command line to an existing script issue the UPDATE SCRIPT command without the LINE= parameter. The appended command line is assigned a line number of five greater than the last command line number in the command line sequence. For example, if your script ends with line 010, the appended command line is assigned a line number of 015.
- **Replacing an existing command**
You can change an existing command line by specifying the LINE= parameter.
- **Adding a command and line number**
You can change an existing script by adding new lines.
- **Deleting a command from a server script**
You can delete an individual command line from a script. When you specify a line number, only the corresponding command line is deleted from the script.

Appending a new command

To append a command line to an existing script issue the UPDATE SCRIPT command without the LINE= parameter. The appended command line is assigned a line number of five greater than the last command line number in the command line sequence. For example, if your script ends with line 010, the appended command line is assigned a line number of 015.

About this task

The following is an example of the QSTATUS script. The script has lines 001, 005, and 010 as follows:

```
001 /* This is the QSTATUS script */
005 QUERY STATUS
010 QUERY PROCESS
```

To append the QUERY SESSION command at the end of the script, issue the following command:

```
update script qstatus "query session"
```

The QUERY SESSION command is assigned a command line number of 015 and the updated script is as follows:

```
001 /* This is the QSTATUS script */
005 QUERY STATUS
010 QUERY PROCESS
015 QUERY SESSION
```

Replacing an existing command

You can change an existing command line by specifying the `LINE=` parameter.

About this task

Line number 010 in the QSTATUS script contains a QUERY PROCESS command. To replace the QUERY PROCESS command with the QUERY STGPOOL command, specify the `LINE=` parameter as follows:

```
update script qstatus "query stgpool" line=10
```

The QSTATUS script is updated to contain the following lines:

```
001 /* This is the QSTATUS script */
005 QUERY STATUS
010 QUERY STGPOOL
015 QUERY SESSION
```

Adding a command and line number

You can change an existing script by adding new lines.

About this task

To add the QUERY NODE command as the new line 007 in the QSTATUS script, issue the following command:

```
update script qstatus "query node" line=7
```

The QSTATUS script is updated to contain the following lines:

```
001 /* This is the QSTATUS script */
005 QUERY STATUS
007 QUERY NODE
010 QUERY STGPOOL
015 QUERY SESSION
```

Deleting a command from a server script

You can delete an individual command line from a script. When you specify a line number, only the corresponding command line is deleted from the script.

About this task

For example, to delete the 007 command line from the QSTATUS script, issue the following command:

```
delete script qstatus line=7
```

Querying a server script to create another server script

You can create more server scripts by querying a script and specifying the `FORMAT=RAW` and `OUTPUTFILE` parameters. You can use the resulting output as input into another script without having to create a script line by line.

About this task

The following example shows how to query the SRTL2 script and direct the output to newscrip.script:

```
query script srtl2 format=raw outputfile=newscrip.script
```

You can then edit the `newscript.script` file with an editor that is available to you on your system. To create a new script by using the edited output from your query, issue:

```
define script srtnew file=newscript.script
```

Running a server script

To process a script, use the `RUN` command. You can run a script that contains substitution variables by specifying them along with the `RUN` command.

About this task

To stop a script that is running, an administrator must halt the server. You cannot cancel a script after it starts by using an IBM Spectrum Protect™ command.

Procedure

- Preview the commands in a script to evaluate the script before you run it. To preview the script without running the commands, enter the `RUN` command with the `PREVIEW=YES` parameter. If the script contains substitution variables, the commands are displayed with the substituted variables.
- Run a script that has no variables by entering the following command: `run qaixc` where `qaixc` is the name of the script.
- Run a script that contains substitution variables by specifying the variable values with the command. Contents of the script:

```
/*-----*/
/* Sample continuation and substitution example */
/* -----*/
SELECT-
$1 FROM-
NODES WHERE-
PLATFORM='$2'
```

To run this script, enter the following command:

```
run qaixc node_name aix
```

Where `node_name` is the value for the `$1` variable and `aix` is the value for the `$2` variable.

Related reference:

[RUN \(Run a server script\)](#)

Administrative client macros

A macro is a file that contains one or more administrative client commands. You can run a macro from the administrative client only in batch or interactive modes. Macros are stored as a file on the administrative client. Macros are not distributed across servers and cannot be scheduled on the server.

Macros can include the following elements:

- Administrative server commands
- Comments
- Continuation characters
- Variables

The name for a macro must follow the naming conventions of the administrative client that is running on your operating system.

In a macro that contains several commands, use the `COMMIT` and `ROLLBACK` commands to control command processing within the macro.

You can include the `MACRO` command within a macro file to call other macros up to 10 levels deep. A macro that is called from the administrative client command line is called a high-level macro. Any macros that are called from within the high-level macro are called *nested* macros.

- Writing commands in a macro
Add administrative commands to a macro. The administrative client ignores any blank lines included in your macro. However, a blank line ends a command that is continued (with a continuation character).

- Writing comments in a macro
Add comments to your macro file to describe the purpose or the commands in it.
- Including continuation characters in a macro
You can use continuation characters in a macro file when you want to run a command that is longer than your screen or window width.
- Including substitution variables in a macro
You can use substitution variables in a macro so that when you run the macro, you can provide values for items such as command parameters. When you use substitution variables, you can use a macro again and again, whenever you need to complete the same task for different objects or with different parameter values.
- Running a macro
Use the MACRO command when you want to run a macro. You can enter the MACRO command in batch or interactive mode.
- Command processing in a macro
When you issue a MACRO command, the server processes all commands in the macro file in order, including commands that are contained in any nested macros. The server commits all commands in a macro after successfully completing processing for the highest-level macro.

Writing commands in a macro

Add administrative commands to a macro. The administrative client ignores any blank lines included in your macro. However, a blank line ends a command that is continued (with a continuation character).

About this task

The following is an example of a macro that is called REG.MAC that registers and grants authority to a new administrator:

```
register admin pease mypasswd -
  contact='david pease, x1234'
grant authority pease -
  classes=policy,storage -
  domains=domain1,domain2 -
  stgpools=stgpool1,stgpool2
```

This example uses continuation characters in the macro file. For more information about continuation characters, see Including continuation characters in a macro.

After you create a macro file, you can update the information that it contains and use it again. You can also copy the macro file. After you have a copy of the macro, you can modify and run the copy.

Writing comments in a macro

Add comments to your macro file to describe the purpose or the commands in it.

About this task

To write a comment:

- Write a slash and an asterisk (/*) to indicate the beginning of the comment.
- Write the comment.
- Write an asterisk and a slash (*/) to indicate the end of the comment.

You can put a comment on a line by itself, or you can put it on a line that contains a command or part of a command.

For example, to use a comment to identify the purpose of a macro, write the following line:

```
/* auth.mac-register new nodes */
```

Or you can write a comment to explain something about a command or part of a command:

```
domain=domain1          /*assign node to domain1 */
```

Comments cannot be nested and cannot span lines. Every line of a comment must contain the comment delimiters.

Including continuation characters in a macro

You can use continuation characters in a macro file when you want to run a command that is longer than your screen or window width.

About this task

Without continuation characters, you can enter up to 256 characters. With continuation characters, you can enter up to 1500 characters. In the MACRO command, values of substitution variables are included in the count of characters.

To use a continuation character, enter a dash or a backslash at the end of the line that you want to continue. With continuation characters, you can continue the following lines of a macro.

Examples

- Continue a command, for example:

```
register admin pease mypasswd -  
contact="david, ext1234"
```

- Continue a list of values by entering a dash or a backslash, with no preceding blank spaces, after the last comma of the list that you enter on the first line. Then, enter the remaining items in the list on the next line with no preceding blank spaces. In the following example a list of storage pool names continues across lines:

```
stgpools=stg1, stg2, stg3, -  
stg4, stg5, stg6
```

- Continue a string of values that are enclosed in quotation marks by entering the first part of the string in quotation marks, followed by a dash or a backslash at the end of the line. Then, enter the remainder of the string on the next line. Enclose the remainder of the string in the same type of quotation marks. The following example shows a string that continues across lines:

```
contact="david pease, bldg. 100, room 2b, san jose,"-  
"ext. 1234, alternate contact-norm pass, ext 2345"
```

The two strings are concatenated with no intervening blanks. You must use only this method to continue a quoted string of values across more than one line.

Including substitution variables in a macro

You can use substitution variables in a macro so that when you run the macro, you can provide values for items such as command parameters. When you use substitution variables, you can use a macro again and again, whenever you need to complete the same task for different objects or with different parameter values.

About this task

A substitution variable consists of a percent sign (%), followed by a unique number that identifies the substitution variable. When you run the file with the MACRO command, you must specify values for the variables.

Restrictions:

- If your system uses the percent sign as a wildcard character, the administrative client interprets a pattern-matching expression in a macro where the percent sign is immediately followed by a digit as a substitution variable.
- You cannot enclose a substitution variable in quotation marks. However, a value that you supply as a substitution for the variable can be a quoted string.

Example

Create a macro that is named AUTH.MAC to register new nodes. The macro has four substitution variables for parameters in the command:

```
/* register new nodes */  
register node %1 %2 - /* userid password */  
contact=%3 - /* 'name, phone number' */  
domain=%4 /* policy domain */
```

When you run the macro, you must enter the values that you want to pass to the server to process the command.

For example, to use the macro to register the node that is named DAVID with a password of DAVIDPW, include a name and phone number as contact information, and assign it to the DOMAIN1 policy domain, enter the following command:

```
macro auth.mac david davidpw "david pease, x1234" domain1
```

Running a macro

Use the MACRO command when you want to run a macro. You can enter the MACRO command in batch or interactive mode.

About this task

If the macro does not contain substitution variables, run the macro by entering the MACRO command with the name of the macro file. For example:

```
macro reg.mac
```

If the macro contains substitution variables, include the values that you want to supply after the name of the macro. Each value is delimited by a space. For example:

```
macro auth.mac pease mypasswd "david pease, x1234" domain1
```

If you enter fewer values than there are substitution variables in the macro, the administrative client replaces the remaining variables with null strings.

If you want to omit one or more values between values, enter a null string (") for each omitted value. For example, if you omit the contact information in the previous example, you must enter:

```
macro auth.mac pease mypasswd "" domain1
```

Related reference:

MACRO (Invoke a macro)

Command processing in a macro

When you issue a MACRO command, the server processes all commands in the macro file in order, including commands that are contained in any nested macros. The server commits all commands in a macro after successfully completing processing for the highest-level macro.

If an error occurs in any command in the macro or in any nested macro, the server stops processing and rolls back any changes that were caused by all previous commands.

If you specify the ITEMCOMMIT option when you enter the DSMADMC command, the server commits each command in a script or a macro individually after successfully completing processing for each command. If an error occurs, the server continues processing and rolls back only the changes caused by the failed command.

You can control precisely when commands are committed with the COMMIT command. If an error occurs while the server is processing the commands in a macro, the server stops processing the macro and rolls back any uncommitted changes. Uncommitted changes are commands that were processed since the last COMMIT command. Make sure that your administrative client session is not running with the ITEMCOMMIT option if you want to control command processing with the COMMIT command.

You can test a macro before you implement it by using the ROLLBACK command. You can enter the commands (except the COMMIT command) you want to issue in the macro, and enter ROLLBACK as the last command. Then, you can run the macro to verify that all the commands process successfully. Any changes to the database caused by the commands are rolled back by the ROLLBACK command. Remember to remove the ROLLBACK command before you make the macro available for actual use. Also, make sure that your administrative client session is not running with the ITEMCOMMIT option if you want to control command processing with the ROLLBACK command.

Tip: Commands that start background processes cannot be rolled back.

If you have a series of commands that process successfully from the command line, but are unsuccessful when issued within a macro, there are probably dependencies between commands. It is possible that a command issued within a macro cannot be processed successfully until a previous command that is issued within the same macro is committed. Either of the following actions allows successful processing of these commands within a macro:

- Insert a COMMIT command before the command dependent on a previous command. For example, if COMMAND C is dependent upon COMMAND B, you would insert a COMMIT command before COMMAND C.

```
command a
command b
commit
command c/
```

- Start the administrative client session by using the ITEMCOMMIT option. This option causes each command within a macro to be committed before the next command is processed.

Related reference:

COMMIT (Control committing of commands in a macro)

ROLLBACK (Rollback uncommitted changes in a macro)

Return codes for use in IBM Spectrum Protect scripts

You can write IBM Spectrum Protect™ scripts that use return codes to determine how script processing proceeds. The return codes can be one of three severities: OK, WARNING, ERROR.

IBM Spectrum Protect scripts use the symbolic return code for processing, not the numeric value. The administrative client displays the numeric values when a command is run. The return codes are shown in the following table.

Table 1. Return codes

Return code	Severity	Numeric value	Description
RC_OK	OK	0	The command completed successfully.
RC_UNKNOWN	ERROR	2	The command is not found; not a known command.
RC_SYNTAX	ERROR	3	The command is valid, but one or more parameters were not specified correctly.
RC_ERROR	ERROR	4	An internal server error prevented the command from successfully completing.
RC_NOMEMORY	ERROR	5	The command could not be completed because of insufficient memory on the server.
RC_NOLOG	ERROR	6	The command could not be completed because of insufficient recovery log space on the server.
RC_NODB	ERROR	7	The command could not be completed because of insufficient database space on the server.
RC_NOSTORAGE	ERROR	8	The command could not be completed because of insufficient storage space on the server.
RC_NOAUTH	ERROR	9	The command failed because the administrator is not authorized to issue the command.
RC_EXISTS	ERROR	10	The command failed because the specified object already exists on the server.
RC_NOTFOUND	WARNING	11	Returned by a QUERY or SQL SELECT command when no objects are found that match specifications.
RC_INUSE	ERROR	12	The command failed because the object to be operated upon was in use.
RC_ISREFERENCED	ERROR	13	The command failed because the object to be operated upon is still referenced by some other server construct.
RC_NOTAVAILABLE	ERROR	14	The command failed because the object to be operated upon is not available.

Return code	Severity	Numeric value	Description
RC_IOERROR	ERROR	15	The command failed because an input/output (I/O) error was encountered on the server.
RC_NOTXN	ERROR	16	The command failed because a database transaction failed on the server.
RC_NOLOCK	ERROR	17	The command failed because a lock conflict was encountered in the server database.
RC_NOTHREAD	ERROR	19	The command could not be completed because of insufficient memory on the server.
RC_LICENSE	ERROR	20	The command failed because the server is not in compliance with licensing.
RC_INVDEST	ERROR	21	The command failed because a destination value was invalid.
RC_IFILEOPEN	ERROR	22	The command failed because an input file that was needed could not be opened.
RC_OFILEOPEN	ERROR	23	The command failed because it could not open a required output file.
RC_OFILEWRITE	ERROR	24	The command failed because it could not successfully write to a required output file.
RC_INVADMIN	ERROR	25	The command failed because the administrator was not defined.
RC_SQLERROR	ERROR	26	An SQL error was encountered during a SELECT statement query.
RC_INVALIDUSE	ERROR	27	The command failed because the command is used in an invalid manner.
RC_NOTABLE	ERROR	28	The command failed because of an unknown SQL table name.
RC_FS_NOTCAP	ERROR	29	The command failed because of incompatible file space name types.
RC_INVALIDADDR	ERROR	30	The command failed because of an incorrect high-level address or low-level address.
RC_INVALIDCG	ERROR	31	The command failed because the management class does not have an archive copy group.
RC_OVERSIZE_VOL	ERROR	32	The command failed because the volume size exceeds the maximum allowed.
RC_DEFVOL_FAIL	ERROR	33	The command failed because volumes cannot be defined in RECLAMATIONTYPE=SNAPLOCK storage pools.
RC_DELVOL_FAIL	ERROR	34	The command failed because volumes cannot be deleted in RECLAMATIONTYPE=SNAPLOCK storage pools.
RC_CANCELED	WARNING	35	The command is canceled.
RC_INVPOLICY	ERROR	36	The command failed because there is an invalid definition in the policy domain.
RC_INVALIDPW	ERROR	37	The command failed because of an invalid password.
RC_UNSUPP_PARM	WARNING	38	The command failed because the command or the parameter is not supported.

Related reference:

DEFINE SCRIPT (Define an IBM Spectrum Protect script)
UPDATE SCRIPT (Update an IBM Spectrum Protect script)
RUN (Run an IBM Spectrum Protect script)

Server documentation in PDF files

Prebuilt PDF files for IBM Spectrum Protect™ documentation are available for you to download.

Tip: Beginning with IBM® Tivoli® Storage Manager Version 7.1.3, the *Administrator's Guide* topics are no longer available in PDF format. Instead, the documentation set is revised to help you complete specific tasks:

- To implement a new data protection solution, see IBM Spectrum Protect data protection solutions. The solution guides provide cookbook-style instructions to help you plan, implement, and manage a solution.
- Alternatively, you can use the IBM Spectrum Protect Blueprints. You can follow the Blueprint procedures to deploy a storage environment, and use the Blueprint scripts to streamline the installation and configuration process. The Blueprints provide the latest hardware and software requirements for small, medium, and large storage environments.
- To administer an *existing* solution, see Configuring servers.

For more information about completing deployment and administration tasks, see the PDF files that are listed in the following table.

Task	Components	Links
Learning about product concepts and solutions	<ul style="list-style-type: none">• Server• Operations Center	Introduction to Data Protection Solutions
Deploying a best practice solution	<ul style="list-style-type: none">• Server• Operations Center	<ul style="list-style-type: none">• Single-Site Disk Solution Guide• Multisite Disk Solution Guide• Tape Solution Guide
Installing components	<ul style="list-style-type: none">• Server• Operations Center	<ul style="list-style-type: none">• AIX®• Linux• Windows
Upgrading components	<ul style="list-style-type: none">• Server	<ul style="list-style-type: none">• AIX• Linux• Windows
Using commands and options	<ul style="list-style-type: none">• Server	<ul style="list-style-type: none">• AIX• Linux• Windows
Using messages and error codes	<ul style="list-style-type: none">• Server	All operating systems

Klienty kopii zapasowych i archiwalnych produktu IBM Spectrum Protect

Klient kopii zapasowych i archiwalnych produktu IBM Spectrum Protect pozwala na zapisywanie kopii plików i katalogów znajdujących się na stacji roboczej lub serwerze plików i przechowywanie ich na serwerze IBM Spectrum Protect. Kopie te można odzyskać w przypadku uszkodzenia lub utraty oryginałów. W zależności od przyczyny tworzenia kopii danych, można tworzyć kopie zapasowe lub archiwalne.

To wydanie nie zawiera zaktualizowanej wersji komponentu klienta kopii zapasowych i archiwalnych. Dokumentacja klienta kopii zapasowych i archiwalnych dostępna jest dla wcześniejszych wydań.

Aplikacyjny interfejs programistyczny

Aplikacyjny interfejs programistyczny (API) produktu IBM Spectrum Protect jest pakietem klienta kopii zapasowych i archiwalnych produktu IBM Spectrum Protect. Używając interfejsu API, można chronić aplikacje biznesowe, takie jak bazy danych, w środowisku IBM Spectrum Protect.

To wydanie nie zawiera zaktualizowanej wersji komponentu API. Dokumentacja interfejsu API dostępna jest dla wcześniejszych wydań.

Wydajność

Na wydajność serwera i klientów ma wpływ wiele czynników, takich jak system operacyjny, sprzęt, konfiguracja sieci, typy pamięci masowej oraz liczba i wielkość plików klientów. Interakcje między tymi czynnikami mogą utrudnić optymalizację wydajności.

To wydanie nie zawiera zaktualizowanej wersji komponentu wydajności. Dokumentację dotyczącą wydajności można znaleźć w serwisie dla wersji 8.1.0.

Rozwiązywanie problemów

Dostępne są procedury lokalizowania i rozwiązywania problemów.

To wydanie nie zawiera zaktualizowanej wersji komponentu rozwiązywania problemów. Dokumentację rozwiązywania problemów można znaleźć w serwisie dla wersji 8.1.0.

Messages, return codes, and error codes

Explanations and suggested actions are available for messages that are issued by IBM Spectrum Protect™ components.

- Introduction to messages
- ANS 0000-9999 messages
- API return codes
- IBM Global Security Kit return codes
The server and client use the IBM Global Security Kit (GSKit) for SSL (Secure Sockets Layer) processing between the server and the backup-archive client. Some messages that are issued for SSL processing include GSKit return codes.
- ANE: Client events logged to the server
- ANR: Server common and platform-specific messages
- I/O error code descriptions in server messages
- Device error codes in the AIX system error log
- [🔗](#) Troubleshooting (V8.1.0 is the most recent publication)

Introduction to messages

Messages, error codes, and return codes are issued by the IBM Spectrum Protect™ server and clients.

Messages and codes can appear on the server console, the administrative client, an operator terminal, the administrative graphical user interface, the backup-archive client, or the hierarchical storage management client (HSM client).

IBM Spectrum Protect provides an activity log to help the administrator track server activity and monitor the system. The activity log contains messages generated by the server, and is stored in the database. The server automatically deletes messages from the activity log after they have passed the specified retention period. Any messages sent to the server console are stored in the activity log. Examples of the types of messages stored in the activity log include:

- When client sessions start or end
- When migration starts or ends
- When backed up files are expired from server storage
- Any output generated from background processes

Some messages have no explanations and are not published. The client can send statistics to the server providing information about a backup or restore. These statistics are informational messages that can be enabled or disabled to the various event logging receivers. These messages are not published.

- IBM Spectrum Protect server and client messages format
- Interpreting return code messages

Related tasks:

- [🔗](#) Using the activity log (V7.1.1)

IBM Spectrum Protect server and client messages format

IBM Spectrum Protect™ server and client messages consist of the following elements:

- A three-letter prefix. Messages have different prefixes to help you identify the IBM Spectrum Protect component that issues the message. Typically, all messages for a component have the same prefix. Sometimes a component issues messages with two or three different prefixes.

For example, backup-archive clients issue messages with the ANS prefix. Backup-archive client events that are logged to the server have the ANE prefix. Server common and server platform-specific messages have the ANR prefix.

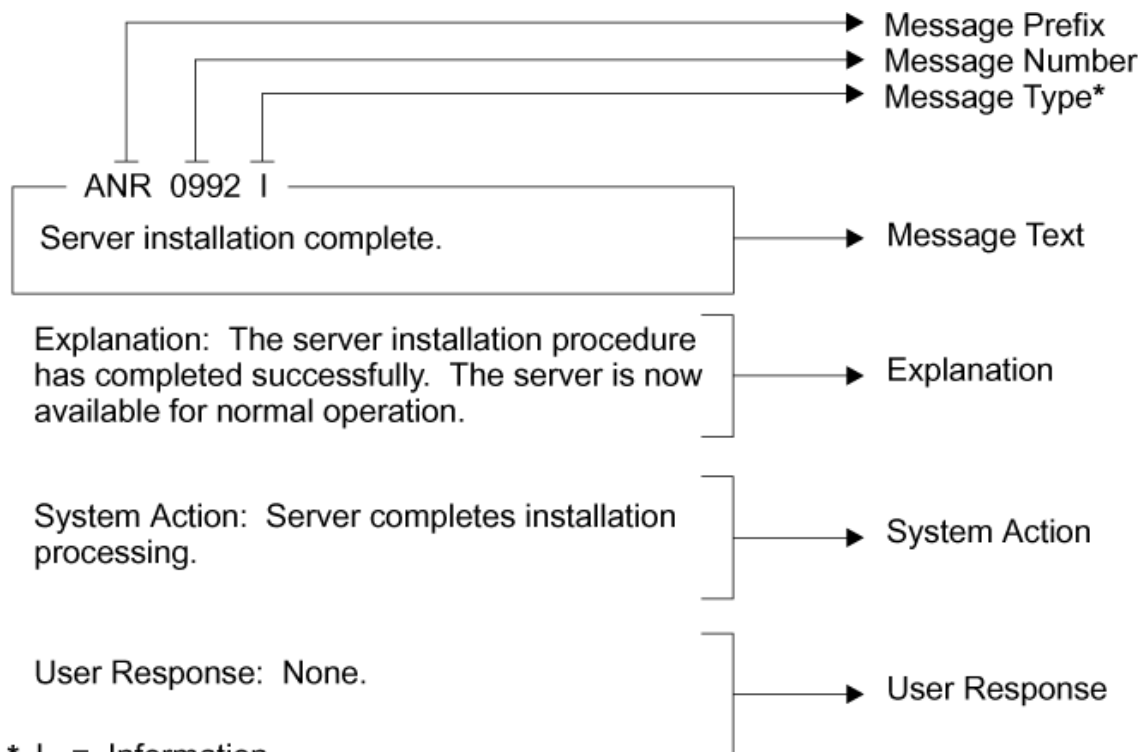
- A numeric message identifier.
- A one-letter severity code. The following codes indicate the severity of the action that generated the message:

Code	Severity	Meaning
S	Severe	The product or a product function cannot continue. User response is required.
E	Error	An error is encountered during processing. Processing might stop. User response might be required.
W	Warning	Processing continues, but problems might occur later as a result of the warning.
I	Information	Processing continues. User response is not necessary.

- Message text that is displayed on screen and written to message logs.
- Explanation, System Action, and User Response texts. These texts elaborate on the message text, and are available in the product messages publications and in the command line help.

The following image presents a typical IBM Spectrum Protect server message.

The callouts identify each element of the message.



- * I = Information
- E = Error
- S = Severe Error
- W = Warning
- K = Kernel message that originates from the hierarchical storage management (HSM) client

Message variables in the message text appear in italics.

Interpreting return code messages

Many different commands can generate the same *return code*. The following examples are illustrations of two different commands issued that result in the same return code; therefore, you must read the *descriptive message* for the command.

In these examples, two different commands yield the same return code, but they also return descriptive messages that are unique to each command. The two commands are `q event standard dddd` and `def vol cstg05 primary`. Both yield a generic message with return code:

```
ANS5102I: Return Code 11.
```

But the first command also yields a descriptive message:

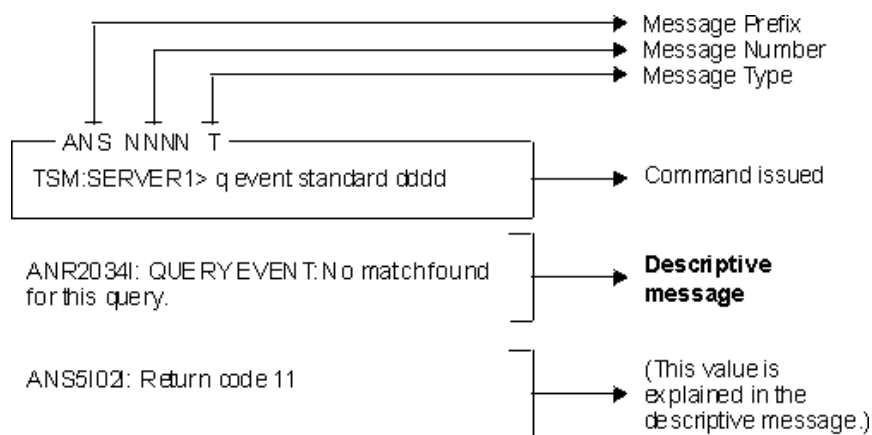
```
ANR2034I: QUERY EVENT: No match found for this query.
```

And the second command also yields a unique, descriptive message:

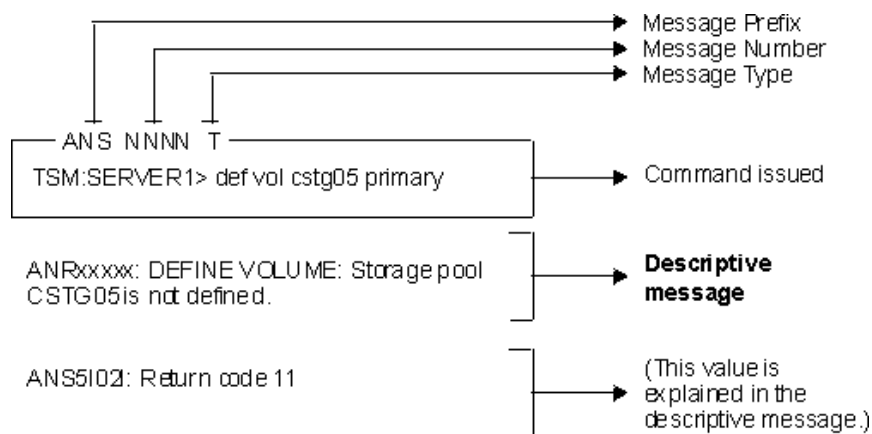
```
ANRxxxx: DEFINE VOLUME: Storage pool CSTG05 is not defined.
```

- Example one for QUERY EVENT command
- Example two for DEFINE VOLUME command

Example one for QUERY EVENT command



Example two for DEFINE VOLUME command



ANE messages

ANE messages are issued by the server. All messages with the ANE prefix are client events logged to the server.

- ANE messages list

ANR messages

ANR messages are issued by the server. Some ANR messages are common to all operating systems, and some are specific to a single operating system.

- ANR messages list

ANS 0000-9999 messages

This release does not include updated ANS-prefix messages. For documentation of ANS-prefix messages, see other releases of IBM Spectrum Protect™.

API return codes

This release does not include updated application programming interface (API) return codes. For documentation of API return codes, see other releases of IBM Spectrum Protect™.

Descriptions of I/O codes in server messages

IBM Spectrum Protect™ messages can contain input/output (I/O) codes. The codes can be operation codes, completion codes, additional sense codes (ASC), and additional sense code qualifier (ASCQ) codes.

Code descriptions are provided for I/O error messages from the IBM Spectrum Protect server for all supported operating systems.

Code

Description

OP

I/O operation that failed. These values can be displayed:

- READ
- WRITE
- FSR (forward space record)
- RSR (reverse space record)
- FSF (forward space file)
- RSF (reverse space file)
- WEOF (write end of file mark)
- OFFL (rewind and unload the tape)
- FLUSH (flush)
- GET_MEDIUM_INFO (get medium information)
- LOCATE (locate)
- QRYLBP (query logical block protection)
- RDBLKID (read block ID)
- SETLBP (set logical block protection)
- SETMODE (set mode)
- REW (rewind)
- SPACEEOD (space end of data)
- TESTREADY (test drive ready)

CC

I/O completion code. This value is returned by the device driver to the server when an error occurs. For a list of completion codes, see Completion code and operation code values overview. For information about tape library system calls and error descriptions for the library I/O control requests, see technote S7002972.

KEY

Byte 2 of the sense bytes from the error. The following lists some definitions:

- 0 = no additional sense bytes available
- 1 = recovered error
- 2 = not ready
- 3 = medium error
- 4 = hardware error

- 5 = incorrect request
- 6 = unit attention (for example, a SCSI bus reset)
- 7 = data protect
- 8 = blank check
- 9 = vendor specific
- A = copy canceled
- B = canceled command
- C = obsolete
- D = volume overflow
- E = miscompare
- F = reserved

ASC/ASCQ

ASC and ASCQ codes are bytes 12 and 13 of the sense bytes. The drive or library reference manual provided with the device contains tables explaining the values of the KEY, ASC, and ASCQ fields. Descriptions of standard ASC and ASCQ codes provides additional information about standard values of ASC and ASCQ codes.

Operating system error codes

When a command fails, the operating system returns an error number. To determine what the error codes mean, take the following action:

- On AIX®, HP-UX, and Solaris, platforms, view the errno.h file in the /usr/include/sys directory. This file provides definitions for error codes.
- On Linux platforms, view the errno-base.h and errno.h files in the /usr/include/asm-generic directory. These files provides definitions for codes.
- On Windows platforms, contact Microsoft Support for help with error messages.
- Completion code and operation code values overview
IBM Spectrum Protect messages can contain device driver completion codes from the device drivers.
- Descriptions of standard ASC and ASCQ codes
Standard ASC and ASCQ codes are described.

Completion code and operation code values overview

IBM Spectrum Protect™ messages can contain device driver completion codes from the device drivers.

- Device drivers completion codes: Common codes
IBM Spectrum Protect device drivers provide completion codes that are common to all device classes.
- Device drivers completion codes: Media changers
IBM Spectrum Protect device drivers provide completion codes that are specific to media changer devices.
- Device drivers completion codes: Tape drives
IBM Spectrum Protect device drivers provide completion codes that are specific to tape drives.

Device drivers completion codes: Common codes

IBM Spectrum Protect™ device drivers provide completion codes that are common to all device classes.

The following table shows common completion code values for IBM Spectrum Protect device drivers. Each entry provides a description for the I/O error message and the recommended action. After completing the recommended action, try the failing operation again.

Table 1. Completion code values common to all device classes

Decimal	Hexadecimal	Description	Recommended action
200	X'C8'	The device indicated a failure condition, but sense data was unavailable.	Try the failing operation again.
201	X'C9'	The device driver failed.	Contact IBM Spectrum Protect Support.
202	X'CA'	The device EEPROM failed.	Test the device. Service the device if necessary.

Decimal	Hexadecimal	Description	Recommended action
203	X'CB'	Manual intervention is required.	Correct the problem on the device. The problem can be a stuck tape, dirty heads, or a jammed library arm.
204	X'CC'	The system recovered from an I/O error; for your information only.	No action necessary.
205	X'CD'	The SCSI adapter failed.	Check for loose cables, bent pins, bad cables, bad SCSI adapters, improper termination, or bad terminators.
206	X'CE'	A general SCSI failure occurred.	Check for loose cables, bent pins, bad cables, bad SCSI adapters, improper termination, or bad terminators.
207	X'CF'	The device cannot perform the requested action.	Ensure that the device is on and ready. Ensure that the drive was defined appropriately with the DEFINE DRIVE command. Ensure that the device class was defined appropriately with the DEFINE DEVCLASS command.
208	X'D0'	The command stopped.	Contact IBM Spectrum Protect Support.
209	X'D1'	A failure is detected in the device microcode.	Check the microcode level of the drive. Contact the drive manufacturer and request the latest level.
210	X'D2'	The device was reset due to device power-up, SCSI bus reset, or manual tape load/eject.	Try the failing operation again.
211	X'D3'	The SCSI bus is busy.	Ensure that the SCSI IDs are correctly assigned to the correct device, and the device is not being accessed by another process.
212	X'D4'	Persistent reservation is not supported on this device.	No action is necessary.
213	X'D5'	A persistent reservation operation failed.	Reset the device and try the operation again. If the problem persists, contact IBM Spectrum Protect Support.

Device drivers completion codes: Media changers

IBM Spectrum Protect™ device drivers provide completion codes that are specific to media changer devices.

The following table shows completion code values for IBM Spectrum Protect device drivers for media changers. Each entry provides a description for the I/O error message and the recommended action. After performing the recommended action, try the failing operation again.

Table 1. Completion code values for media changers

Decimal	Hexadecimal	Description	Recommended action
300	X'12C'	Cartridge entry/exit error	Check the entry/exit ports for a jammed volume.
301	X'12D'	Cartridge load failure	Check the drive for jammed volumes. On AIX®, display the errpt to check for hardware errors.
302	X'12E'	Cartridge in failed drive	Check the drive for jammed volumes. On AIX, display the errpt to check for hardware errors.

Decimal	Hexadecimal	Description	Recommended action
303	X'12F'	Carousel not loaded	Ensure that the carousel is correctly in place and the door is shut.
304	X'130'	Changer failure	On AIX, display the errpt to check for hardware errors.
305	X'131'	Drive failure	Ensure that the heads are clean. On AIX, display the errpt to check for hardware errors.
306	X'132'	Drive or media failure	Ensure that the heads are clean. On AIX, display the errpt to check for hardware errors.
307	X'133'	Entry/exit failure	Check the tape library interface for hardware errors. If there are no errors, contact IBM Spectrum Protect Support.
308	X'134'	Entry/exit port not present	Check the tape library interface for hardware errors. If there are no errors, contact IBM Spectrum Protect Support.
309	X'135'	Library audit error	Ensure that there are no jammed volumes. It is possible that the library audit is failing due to hardware errors. On AIX, display the errpt to check for hardware errors.
310	X'136'	Library full	Check for jammed volumes. Ensure that the volumes are not rearranged. If the library is not full, start the AUDIT LIBRARY command.
311	X'137'	Media export	Check the tape library interface for hardware errors. If there are no errors, contact IBM Spectrum Protect Support.
312	X'138'	Slot failure	Ensure that nothing is jammed in the slot.
313	X'139'	Slot or media failure	Ensure that the volume is not jammed in the slot and that the volumes are not rearranged. If the problem persists, start the AUDIT LIBRARY command.
314	X'13A'	The source slot or drive was empty in an attempt to move a volume	Ensure that the volumes are not rearranged. If the problem persists, start the AUDIT LIBRARY command.
315	X'13B'	The destination slot or drive was full in an attempt to move a volume	Ensure that the volumes are not rearranged, or that a volume is not stuck in the drive. If problem persists, start the AUDIT LIBRARY command.
316	X'13C'	Cleaner cartridge installed	Contact IBM Spectrum Protect support.
317	X'13D'	Media not ejected	Ensure that the volumes are not rearranged, or that a volume is not stuck in the drive. If problem persists, start the AUDIT LIBRARY command.
318	X'13E'	I/O port not configured	Contact IBM Spectrum Protect Support.
319	X'13F'	First destination empty	Ensure that the volumes are not rearranged. If problem persists, start the AUDIT LIBRARY command.
320	X'140'	No inventory information	Start the AUDIT LIBRARY command.

Decimal	Hexadecimal	Description	Recommended action
321	X'141'	Read element status mismatch	Ensure that host bus adapter drivers and firmware are at current levels. Check the tape library interface for hardware errors. If there are no errors, contact IBM Spectrum Protect Support.
322	X'142'	Initialize range failed	Check the tape library interface for hardware errors. If there are no errors, contact IBM Spectrum Protect Support.

Device drivers completion codes: Tape drives

IBM Spectrum Protect™ device drivers provide completion codes that are specific to tape drives.

The following table shows completion code values for IBM Spectrum Protect device drivers for tape drives. Each entry provides a description for the I/O error message and the recommended action. After trying the recommended action, try the failing operation again.

Table 1. Completion code values for tape drives

Decimal	Hexadecimal	Description	Recommended action
400	X'190'	Physical end of media encountered	Ensure that the heads are clean on the drive.
401	X'191'	End of data detected	Contact IBM Spectrum Protect Support.
402	X'192'	Media corrupted	Ensure that the heads are clean. Ensure that the media is not physically damaged and has not reached the end of life as specified by the media manufacturer.
403	X'193'	Media failure	Ensure that the heads are clean. Ensure that the media is not physically damaged and has not reached the end of life as specified by the media manufacturer.
404	X'194'	Media incompatibility	Ensure that the correct length and type of media is being used.
406	X'196'	Sector that is requested is invalid	Internal server error. Contact IBM Spectrum Protect Support.
407	X'197'	Write protect	Ensure that the volume is not write protected.
408	X'198'	Clean the media and the drive	Clean the drive heads with a cleaning cartridge.
409	X'199'	Media fault	Ensure that the heads are clean. Ensure that the media is not physically damaged and has not reached the end of life as specified by the media manufacturer.
410	X'19A'	Cleaning complete	Try the failing operation again.
411	X'19B'	Logical end of media encountered	Contact IBM Spectrum Protect Support.
412	X'19C'	Media not present in drive	Ensure that the media is correctly positioned in the drive. If problem persists, start the AUDIT LIBRARY command.
413	X'19D'	Encountered the beginning of the media	Contact IBM Spectrum Protect Support.
414	X'19E'	Erase failure	Clean the drive heads.
415	X'19F'	Attempted to overwrite written WORM media	Internal server error. Contact IBM Spectrum Protect Support.

Decimal	Hexadecimal	Description	Recommended action
416	X'1A0'	An incorrect length block was read.	Ensure that the heads are clean. On AIX®, display the errpt to check for hardware errors.
417	X'1A1'	Open read only	Contact IBM Spectrum Protect Support.
418	X'1A2'	Open write only	Contact IBM Spectrum Protect Support.
419	X'1A2'	Media scan failed	Clean the drive and media.
420	X'1A4'	Logical write protect	Ensure that the heads are clean. Check operating system error logs for hardware errors. Verify that the write protect tab is off. Turn off SAN tape acceleration or set CHECKTAPEPOS to OFF or TSMonly.
422	X'1A6'	Cleaning is required	Clean the tape drive.
423	X'1A7'	Media error	Check operating system error logs for hardware errors. Check for bad media.
424	X'1A8'	Encryption-related error occurred	Check your encryption setting on your device class and tape drive.
425	X'1A9'	Decryption-related error occurred	Check your encryption setting on your device class and tape drive.
425	X'1AA'	An external, encryption-related error occurred	Check the encryption setting on your device class and tape drive.
426	X'1AB'	A CRC mismatch occurred	Ensure that the media has not reached the end of life as specified by the media manufacturer. Try the operation again.

Descriptions of standard ASC and ASCQ codes

Standard ASC and ASCQ codes are described.

The ASC and ASCQ codes are bytes 12 and 13 for SCSI-2 devices. On Windows systems, these codes are displayed in the Windows Event Log, but the information is in different bytes.

See server message ANR8300E or ANR8302E for the recommended action.

The following table provides standard descriptions for some ASC and ASCQ codes. Each value has a prefix of 0x, which indicates that it is a hexadecimal constant. Note that descriptions vary among devices. For an accurate description of ASC and ASCQ codes for any device, see the documentation that comes with the device.

Table 1. Descriptions of standard ASC and ASCQ codes

ASC	ASCQ	Description
0x00	0x00	No additional sense
0x00	0x01	Filemark detected
0x00	0x02	End-of-medium detected
0x00	0x03	Setmark detected
0x00	0x04	Beginning of medium
0x00	0x05	End of data
0x00	0x06	I/O process terminated
0x02	0x00	No seek complete
0x03	0x00	Device write fault
0x03	0x01	No write current

ASC	ASCQ	Description
0x03	0x02	Excessive write errors
0x04	0x00	Logical unit not ready
0x04	0x01	Becoming ready
0x04	0x02	Not ready, initializing command required
0x04	0x03	Not ready, manual intervention required
0x04	0x04	Not ready, formatting
0x05	0x00	No response to select
0x06	0x00	No reference position found
0x07	0x00	Multiple devices selected
0x08	0x00	Communication failure
0x08	0x01	Communication timeout
0x08	0x02	Communication parity error
0x09	0x00	Track following error
0x0A	0x00	Error log overflow
0x0C	0x00	Write error
0x11	0x00	Unrecovered read error
0x11	0x01	Read retries exhausted
0x11	0x02	Error too long to correct
0x11	0x03	Multiple read errors
0x11	0x08	Incomplete block read
0x11	0x09	No gap found
0x11	0x0A	Miscorrected error
0x14	0x00	Recorded entity not found
0x14	0x01	Record not found
0x14	0x02	Filemark/setmark not found
0x14	0x03	End-of-data not found
0x14	0x04	Block sequence error
0x15	0x00	Random positioning error
0x15	0x01	Mechanical positioning error
0x15	0x02	Read positioning error
0x17	0x00	No error correction applied
0x17	0x01	Recovered with retries
0x17	0x02	Recovered with positive head offset
0x17	0x03	Recovered with negative head offset
0x18	0x00	ECC applied
0x1A	0x00	Parameter list length error
0x1B	0x00	Synchronous data transfer error
0x20	0x00	Invalid operation code
0x21	0x00	Block out of range
0x21	0x01	Invalid element address

ASC	ASCQ	Description
0x24	0x00	Invalid field in CDB
0x25	0x00	LUN not supported
0x26	00	Invalid field in parameter list
0x26	0x01	Parameter not supported
0x26	0x02	Parameter value invalid
0x26	0x03	Threshold parameters not supported
0x27	0x00	Write protected
0x28	0x00	Not-ready to ready
0x28	0x01	Import/export element accessed
0x29	0x00	Power-on, reset, bus reset
0x2A	0x00	Parameters changed
0x2A	0x01	Mode parameters changed
0x2A	0x02	Log parameters changed
0x2B	0x00	Copy cannot run
0x2C	0x00	Command sequence error
0x2D	0x00	Overwrite error on update
0x2F	0x00	Command cleared by initiator
0x30	0x00	Incompatible media
0x30	0x01	Media unknown format
0x30	0x02	Media incompatible format
0x30	0x03	Cleaning cartridge installed
0x31	0x00	Media format corrupted
0x33	0x00	Tape length error
0x37	0x00	Rounded parameter
0x39	0x00	Saving parameters not supported
0x3A	0x00	Medium not present
0x3B	0x00	Sequential positioning error
0x3B	0x01	Positioning error at BOT
0x3B	0x02	Positioning error at EOT
0x3B	0x08	Reposition error
0x3B	0x0D	Medium destination element full
0x3B	0x0E	Medium source element empty
0x3D	0x00	Invalid bits in message
0x3E	0x00	LUN not self-configured
0x3F	0x00	Operating conditions changed
0x3F	0x01	Microcode changed
0x3F	0x02	Changed operating definition
0x3F	0x03	Inquiry data changed
0x3F	0x0E	Reported LUNs data changed
0x43	0x00	Message error

ASC	ASCQ	Description
0x44	0x00	Internal target failure
0x45	0x00	Select/reselect failure
0x46	0x00	Unsuccessful soft reset
0x47	0x00	SCSI parity error
0x48	0x00	Initiator detected message received
0x49	0x00	Invalid message error
0x4A	0x00	Command phase error
0x4B	0x00	Data phase error
0x4C	0x00	LUN failed self-configuration
0x4E	0x00	Overlapped commands attempt
0x50	0x00	Write append error
0x50	0x01	Write append position error
0x50	0x02	Position error (timing)
0x51	0x00	Erase failure
0x52	0x00	Cartridge fault
0x53	0x00	Load/media eject failed
0x53	0x01	Unload tape failure
0x53	0x02	Media removal prevented
0x5A	0x00	Operator state changed
0x5A	0x01	Operator media removal
0x5A	0x02	Operator write protect
0x5A	0x03	Operator write permit
0x5B	0x00	Log exception
0x5B	0x01	Threshold condition met
0x5B	0x02	Log counter at maximum
0x5B	0x03	Log list codes exhausted

- ASC and ASCQ codes in the Windows Event Log
ASC and ASCQ codes are displayed in the Windows Event Log.

Device error codes in the AIX system error log

Some device error codes are logged in the AIX® system error log.

ADSM_DD_LOG1 (0xAC3AB953)
DEVICE DRIVER SOFTWARE ERROR

This error is logged by the IBM Spectrum Protect™ device driver when a problem is suspected in the IBM Spectrum Protect device driver software. If the IBM Spectrum Protect device driver issues a SCSI I/O command with an illegal operation code, the command fails and the error is logged with this identifier. Report this error immediately to IBM Spectrum Protect Support.

Detail Data: Sense Data

The sense data contains information that can determine the cause of the error. Report all data in the error entry to IBM Spectrum Protect Support.

ADSM_DD_LOG2 (0x5680E405)

HARDWARE/COMMAND-ABORTED ERROR

This error is logged by the IBM Spectrum Protect device driver when the device reports a hardware error or stop-command error in response to a SCSI I/O command.

Detail Data: Sense Data

The sense data contains information that can determine which hardware component failed and why. To interpret the sense data for a particular device, refer to the SCSI specification manual for the device.

ADSM_DD_LOG3 (0x461B41DE) MEDIA ERROR

This error is logged by the IBM Spectrum Protect device driver when a SCSI I/O command fails because of corrupted or incompatible media, or because a drive requires cleaning.

Detail Data: Sense Data

The sense data contains information that can determine the cause of the error. To interpret the sense data for a particular device, refer to the SCSI specification manual for the device.

ADSM_DD_LOG4 (0x4225DB66) TARGET DEVICE GOT UNIT ATTENTION

This error is logged by the IBM Spectrum Protect device driver after receiving certain UNIT ATTENTION notifications from a device. UNIT ATTENTIONs are informational and usually indicate that some state of the device changed. For example, this error would be logged if the door of a library device was opened and then closed. Logging this event indicates that the activity occurred and that the library inventory might be changed.

Detail Data: Sense Data

The sense data contains information that describes the reason for the UNIT ATTENTION. To interpret the sense data for a particular device, see the SCSI specification manual for the device.

ADSM_DD_LOG5 (0xDAC55CE5) PERMANENT UNKNOWN ERROR

This error is logged by the IBM Spectrum Protect device driver after receiving an unknown error from a device in response to a SCSI I/O command. If the error persists, report it to IBM Spectrum Protect support personnel.

Detail Data: Sense Data

The sense data consists of information that can determine the cause of the error. Report all data in the error entry to IBM Spectrum Protect Support.

ADSM_DD_LOG6 (0xBC539B26) WARNING OR INFORMATIONAL MESSAGE FOR TARGET DEVICE

This error is logged by the IBM Spectrum Protect device driver after receiving a warning or informational message from a device in response to a SCSI I/O command. These warning or informational messages might not be an indication of a problem. They could be an indication that cleaning is completed, that the cleaning cartridge is inserted, or something similar. If the message persists, report it to IBM Spectrum Protect Support.

Detail Data: Sense Data

The sense data consists of information that can determine the reason for the message. Report all data in the entry to IBM Spectrum Protect Support.

IBM Global Security Kit return codes

The server and client use the IBM Global Security Kit (GSKit) for SSL (Secure Sockets Layer) processing between the server and the backup-archive client. Some messages that are issued for SSL processing include GSKit return codes.

GSKit is automatically installed or updated during IBM Spectrum Protect™ installation and provides the following libraries:

- GSKit SSL

- GSKit Key Management API
- IBM Crypto for C (ICC)

The `tsmdiag` utility reports the GSKit level that is installed on your system, or you can use one of the following methods:

- For Windows, issue the following commands:

```
regedit /e gskitinfo.txt "HKEY_LOCAL_MACHINE\software\ibm\gsk8\"
notepad gskitinfo.txt
```

CAUTION:

You can damage the system registry if you use `regedit` incorrectly.

- For the 64-bit AIX® server, issue the following command from the command line: `gsk8ver_64`

See Table 1 for the GSKit SSL return codes.

The server uses the GSKit Key Management API to automatically create the key management database and server private and public keys. Some messages that are issued for this processing might include GSKit Key Management return codes. See Table 2 for the key management return codes.

Table 1. IBM Global Security Kit SSL general return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completes successfully. Issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful <code>open()</code> function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The dynamic link library (DLL) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to complete the operation.
0x00000005	5	GSK_INVALID_STATE	The handle is not in a valid state for operation, such as completing an <code>init()</code> operation on a handle twice.
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label is not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate is not received from the partner.
0x00000008	8	GSK_ERROR_CERT_VALIDATION	Certificate validation error.
0x00000009	9	GSK_ERROR_CRYPTO	Error processing cryptography.
0x0000000a	10	GSK_ERROR_ASN	Error validating ASN fields in certificate.
0x0000000b	11	GSK_ERROR_LDAP	Error connecting to user registry.
0x0000000c	12	GSK_ERROR_UNKNOWN_ERROR	Internal error. Report this error to IBM Software Support.
0x0000000d	13	GSK_INVALID_PARAMETER	Invalid parameter.
0x0000000e	14	GSK_ERROR_UNEXPECTED_INT_EXCEPTION	Invalid parameter. Report this error to IBM Software Support.
0x00000065	101	GSK_OPEN_CIPHER_ERROR	Internal error. Report this error to IBM Software Support.
0x00000066	102	GSK_KEYFILE_IO_ERROR	I/O error reading the key file.
0x00000067	103	GSK_KEYFILE_INVALID_FORMAT	The key file does not have a valid internal format. Recreate the key file.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000068	104	GSK_KEYFILE_DUPLICATE_KEY	The key file has two entries with the same key.
0x00000069	105	GSK_KEYFILE_DUPLICATE_LABEL	The key file has two entries with the same label.
0x0000006a	106	GSK_BAD_FORMAT_OR_INVALID_PASSWORD	The key file password is used as an integrity check. Either the key file is corrupted or the password ID is incorrect.
0x0000006b	107	GSK_KEYFILE_CERT_EXPIRED	The default key in the key file has an expired certificate.
0x0000006c	108	GSK_ERROR_LOAD_GSKLIB	An error occurred loading one of the GSK dynamic link libraries. Check that GSK was installed correctly.
0x0000006d	109	GSK_PENDING_CLOSE_ERROR	Indicates that a connection is trying to be made in a GSK environment after the GSK_ENVIRONMENT_CLOSE_OPTIONS was set to GSK_DELAYED_ENVIRONMENT_CLOSE and gsk_environment_close() function was called.
0x000000c9	201	GSK_NO_KEYFILE_PASSWORD	Both the password and the stash-file name were not specified. The key file is not initialized.
0x000000ca	202	GSK_KEYRING_OPEN_ERROR	Unable to open the key file. Either the path was specified incorrectly or the file permissions did not allow the file to be opened.
0x000000cb	203	GSK_RSA_TEMP_KEY_PAIR	Unable to generate a temporary key pair. Report this error to IBM Software Support.
0x000000cc	204	GSK_ERROR_LDAP_NO_SUCH_OBJECT	A user name object was specified that is not found.
0x000000cd	205	GSK_ERROR_LDAP_INVALID_CREDENTIALS	A password that is used for an LDAP (lightweight directory access protocol) query is not correct.
0x000000ce	206	GSK_ERROR_BAD_INDEX	An index into the Fail Over list of LDAP servers was not correct.
0x000000cf	207	GSK_ERROR_FIPS_NOT_SUPPORTED	This installation of GSKit does not support FIPS mode of operation.
0x0000012d	301	GSK_CLOSE_FAILED	Indicates that the GSK environment close request was not properly managed. Cause is most likely due to a gsk_secure_socket*() command that is attempted after a gsk_close_environment() call.
0x00000191	401	GSK_ERROR_BAD_DATE	The system date was not set to a valid value.
0x00000192	402	GSK_ERROR_NO_CIPHERS	The SSLv2 and the SSLv3 are not enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from the partner.
0x00000194	404	GSK_ERROR_BAD_CERTIFICATE	The received certificate was formatted incorrectly.
0x00000195	405	GSK_ERROR_UNSUPPORTED_CERTIFICATE_TYPE	The received certificate type was not supported.
0x00000196	406	GSK_ERROR_IO	An I/O error occurred on a data read or write operation.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000197	407	GSK_ERROR_BAD_KEYFILE_LABEL	The specified label in the key file is not found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file cannot be used. The key file also might be corrupt.
0x00000199	409	GSK_ERROR_BAD_KEY_LEN_FOR_EXPORT	In a restricted cryptography environment, the key size is too long to be supported.
0x0000019a	410	GSK_ERROR_BAD_MESSAGE	An incorrectly formatted SSL message was received from the partner.
0x0000019b	411	GSK_ERROR_BAD_MAC	The message authentication code (MAC) was not successfully verified.
0x0000019c	412	GSK_ERROR_UNSUPPORTED	Unsupported SSL protocol or unsupported certificate type.
0x0000019d	413	GSK_ERROR_BAD_CERT_SIG	The received certificate contained an incorrect signature.
0x0000019e	414	GSK_ERROR_BAD_CERT	Incorrectly formatted certificate is received from the partner.
0x0000019f	415	GSK_ERROR_BAD_PEER	Did not receive a valid SSL protocol from the partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.
0x000001a1	417	GSK_ERROR_SELF_SIGNED	The self-signed certificate is not valid.
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The <code>read()</code> failed. Report this error to IBM Software Support.
0x000001a3	419	GSK_ERROR_NO_WRITE_FUNCTION	The <code>write()</code> failed. Report this error to IBM Software Support.
0x000001a4	420	GSK_ERROR_SOCKET_CLOSED	The partner closed the socket before the protocol completed.
0x000001a5	421	GSK_ERROR_BAD_V2_CIPHER	The specified V2 cipher is not valid.
0x000001a6	422	GSK_ERROR_BAD_V3_CIPHER	The specified V3 cipher is not valid.
0x000001a7	423	GSK_ERROR_BAD_SEC_TYPE	Report this error to IBM Software Support.
0x000001a8	424	GSK_ERROR_BAD_SEC_TYPE_COMBINATION	Report this error to IBM Software Support.
0x000001a9	425	GSK_ERROR_HANDLE_CREATION_FAILED	The handle cannot be created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Not able to access the specified user registry when a certificate is being validated.
0x000001ac	428	GSK_ERROR_NO_PRIVATE_KEY	The specified key did not contain a private key.
0x000001ad	429	GSK_ERROR_PKCS11_LIBRARY_NOTLOADED	A failed attempt was made to load the specified PKCS11 shared library.
0x000001ae	430	GSK_ERROR_PKCS11_TOKEN_LABELMISMATCH	The PKCS #11 driver failed to find the token that is specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is not valid.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001b2	434	GSK_CSP_OPEN_ERROR	Cannot open the hardware-based cryptographic service provider. Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CONFLICTING_ATTRIBUTE_SETTING	Attribute setting conflict between PKCS11, CMS key database, and Microsoft Crypto API.
0x000001b4	436	GSK_UNSUPPORTED_PLATFORM	The requested function is not supported on the platform that the application is running. For example, the Microsoft Crypto API is not supported on platforms other than Windows 2000.
0x000001b6	438	GSK_ERROR_INCORRECT_SESSION_TYPE	Incorrect value is returned from the reset session type callback function. Only GSKit <code>gsk_sever_session</code> , <code>gsk_sever_session_with_cl_auth</code> , or <code>gsk_sever_session_with_cl_auth_crit</code> is allowed.
0x000001f5	501	GSK_INVALID_BUFFER_SIZE	The buffer size is negative or zero.
0x000001f6	502	GSK_WOULD_BLOCK	Used with nonblocking I/O. Refer to the nonblocking section for usage.
0x00000259	601	GSK_ERROR_NOT_SSLV3	SSLv3 is required for <code>reset_cipher()</code> , and the connection uses SSLv2.
0x0000025a	602	GSK_MISC_INVALID_ID	A valid ID was not specified for the <code>gsk_secure_soc_misc()</code> function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call does not have a valid ID. This issue might also be caused by specifying an environment handle when a handle for an SSL connection should be used.
0x000002be	702	GSK_ATTRIBUTE_INVALID_LENGTH	The attribute has a negative length, which is not valid.
0x000002bf	703	GSK_ATTRIBUTE_INVALID_ENUMERATION	The enumeration value is not valid for the specified enumeration type.
0x000002c0	704	GSK_ATTRIBUTE_INVALID_SID_CACHE	A parameter list that is not valid for replacing the SID cache routines.
0x000002c1	705	GSK_ATTRIBUTE_INVALID_NUMERIC_VALUE	When a numeric attribute is set, the specified value is not valid for the specific attribute that is being set.
0x000002c2	706	GSK_CONFLICTING_VALIDATION_SETTING	Conflicting parameters were set for additional certificate validation.
0x000002c3	707	GSK_AES_UNSUPPORTED	The AES cryptographic algorithm is not supported.
0x000002c4	708	GSK_PEERID_LENGTH_ERROR	The PEERID does not have the correct length.
0x000002c5	709	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_OFF	The particular cipher is not allowed when FIPS mode of operation is off.
0x000002c6	710	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_ON	No approved FIPS ciphers are selected in FIPS mode of operation.
0x00000641	1601	GSK_TRACE_STARTED	The trace started successfully.
0x00000642	1602	GSK_TRACE_STOPPED	The trace stopped successfully.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000643	1603	GSK_TRACE_NOT_STARTED	No trace file was previously started, so it cannot be stopped.
0x00000644	1604	GSK_TRACE_ALREADY_STARTED	Trace file is started, so it cannot be restarted.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid full path file name.

Table 2. IBM Global Security Kit key management return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completes successfully. This message is issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful <code>open()</code> function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The DLL (dynamic link library) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to complete the operation.
0x00000005	5	GSK_INVALID_STATE	The handle is in an incorrect state for operation, such as completing an <code>init()</code> operation on a handle twice.
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label is not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate is not received from the partner.
0x00000008	8	GSK_ERROR_CERT_VALIDATION	Certificate validation error.
0x00000009	9	GSK_ERROR_CRYPTO	Error processing cryptography.
0x0000000a	10	GSK_ERROR_ASN	Error validating ASN fields in certificate.
0x0000000b	11	GSK_ERROR_LDAP	Error connecting to user registry.
0x0000000c	12	GSK_ERROR_UNKNOWN_ERROR	Internal error. Report this error to IBM Software Support.
0x00000065	101	GSK_OPEN_CIPHER_ERROR	Internal error. Report this error to IBM Software Support.
0x00000066	102	GSK_KEYFILE_IO_ERROR	I/O error reading the key file.
0x00000067	103	GSK_KEYFILE_INVALID_FORMAT	The key file has an internal format that is not valid. Recreate key file.
0x00000068	104	GSK_KEYFILE_DUPLICATE_KEY	The key file has two entries with the same key.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000069	105	GSK_KEYFILE_DUPLICATE_LABEL	The key file has two entries with the same label.
0x0000006a	106	GSK_BAD_FORMAT_OR_INVALID_PASSWORD	The key file password is used as an integrity check. Either the key file is corrupted or the password ID is incorrect.
0x0000006b	107	GSK_KEYFILE_CERT_EXPIRED	The default key in the key file has an expired certificate.
0x0000006c	108	GSK_ERROR_LOAD_GSKLIB	An error occurred while one of the GSK dynamic link libraries is loaded. Check GSK was installed correctly.
0x0000006d	109	GSK_PENDING_CLOSE_ERROR	This message indicates that a connection is trying to be made in a GSK environment after the GSK_ENVIRONMENT_CLOSE_OPTIONS was set to GSK_DELAYED_ENVIRONMENT_CLOSE and gsk_environment_close() function was called.
0x000000c9	201	GSK_NO_KEYFILE_PASSWORD	Both the password and the stash-file name were not specified, so the key file is not initialized.
0x000000ca	202	GSK_KEYRING_OPEN_ERROR	Unable to open the key file. Either the path was specified incorrectly or the file permissions did not allow the file to be opened.
0x000000cb	203	GSK_RSA_TEMP_KEY_PAIR	Unable to generate a temporary key pair. Report this error to IBM Software Support.
0x000000cc	204	GSK_ERROR_LDAP_NO_SUCH_OBJECT	A user name object was specified that is not found.
0x000000cd	205	GSK_ERROR_LDAP_INVALID_CREDENTIALS	A Password that is used for an LDAP query is not correct.
0x000000ce	206	GSK_ERROR_BAD_INDEX	An index into the Fail Over list of LDAP servers was not correct.
0x000000cf	207	GSK_ERROR_FIPS_NOT_SUPPORTED	This installation of GSKit does not support FIPS mode of operation.
0x0000012d	301	GSK_CLOSE_FAILED	Indicates that the GSK environment close request was not properly managed. Cause is most likely due to attempting a gsk_secure_socket*() command after a gsk_close_environment() call.
0x00000191	401	GSK_ERROR_BAD_DATE	The system date was set to a value that is not valid.
0x00000192	402	GSK_ERROR_NO_CIPHERS	SSLv2 and SSLv3 are not enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from the partner.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000194	404	GSK_ERROR_BAD_CERTIFICATE	The received certificate was formatted incorrectly.
0x00000195	405	GSK_ERROR_UNSUPPORTED_CERTIFICATE_TYPE	The received certificate type was not supported.
0x00000196	406	GSK_ERROR_IO	An I/O error occurred on a data read-or-write operation.
0x00000197	407	GSK_ERROR_BAD_KEYFILE_LABEL	The specified label in the key file is not found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file cannot be used. The key file might also be corrupt.
0x00000199	409	GSK_ERROR_BAD_KEY_LEN_FOR_EXPORT	In a restricted cryptography environment, the key size is too long to be supported.
0x0000019a	410	GSK_ERROR_BAD_MESSAGE	An incorrectly formatted SSL message was received from the partner.
0x0000019b	411	GSK_ERROR_BAD_MAC	The MAC was not successfully verified.
0x0000019c	412	GSK_ERROR_UNSUPPORTED	Unsupported SSL protocol or unsupported certificate type.
0x0000019d	413	GSK_ERROR_BAD_CERT_SIG	The received certificate contained an incorrect signature.
0x0000019e	414	GSK_ERROR_BAD_CERT	Incorrectly formatted certificate is received from the partner.
0x0000019f	415	GSK_ERROR_BAD_PEER	An SSL protocol that is not valid is received from the partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.
0x000001a1	417	GSK_ERROR_SELF_SIGNED	The self-signed certificate is not valid.
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The read() failed. Report this error to IBM Software Support.
0x000001a3	419	GSK_ERROR_NO_WRITE_FUNCTION	The write() failed. Report this error to IBM Software Support.
0x000001a4	420	GSK_ERROR_SOCKET_CLOSED	The partner closed the socket before the protocol completed.
0x000001a5	421	GSK_ERROR_BAD_V2_CIPHER	The specified V2 cipher is not valid.
0x000001a6	422	GSK_ERROR_BAD_V3_CIPHER	The specified V3 cipher is not valid.
0x000001a7	423	GSK_ERROR_BAD_SEC_TYPE	Report this error to IBM Software Support.
0x000001a8	424	GSK_ERROR_BAD_SEC_TYPE_COMBINATION	Report this error to IBM Software Support.
0x000001a9	425	GSK_ERROR_HANDLE_CREATION_FAILED	The handle is not created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Unable to access the specified user registry when a certificate is being validated
0x000001ac	428	GSK_ERROR_NO_PRIVATE_KEY	The specified key did not contain a private key.
0x000001ad	429	GSK_ERROR_PKCS11_LIBRARY_NOTLOADED	A failed attempt was made to load the specified PKCS11 shared library.
0x000001ae	430	GSK_ERROR_PKCS11_TOKEN_LABELMISMATCH	The PKCS #11 driver failed to find the token that is specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is incorrect.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.
0x000001b2	434	GSK_CSP_OPEN_ERROR	Could not open the hardware-based cryptographic service provider (CSP). Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CSP_OPEN_ERROR	Some conflicting attributes for SSL operation were defined.
0x000001b4	436	GSK_CSP_OPEN_ERROR	The Microsoft Crypto API is only supported on Microsoft Windows 2000 with Service Pack 2 applied.
0x000001b5	437	GSK_CSP_OPEN_ERROR	System is running in IPv6 mode without setting a PEERID.
0x000001f5	501	GSK_INVALID_BUFFER_SIZE	The buffer size is negative or zero.
0x000001f6	502	GSK_WOULD_BLOCK	Used with nonblocking I/O. Refer to the nonblocking section for usage.
0x00000259	601	GSK_ERROR_NOT_SSLV3	SSLv3 is required for reset_cipher(), and the connection uses SSLv2.
0x0000025a	602	GSK_MISC_INVALID_ID	An ID that is not valid was specified for the gsk_secure_soc_misc() function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call has an ID that is not valid. This issue might also be caused by specifying an environment handle when a handle for an SSL connection should be used.
0x000002be	702	GSK_ATTRIBUTE_INVALID_LENGTH	The attribute has a negative length, which is not valid.
0x000002bf	703	GSK_ATTRIBUTE_INVALID_ENUMERATION	The enumeration value is not valid for the specified enumeration type.

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000002c0	704	GSK_ATTRIBUTE_INVALID_SID_CACHE	A parameter list that is not valid for replacing the SID cache routines.
0x000002c1	705	GSK_ATTRIBUTE_INVALID_NUMERIC_VALUE	When a numeric attribute is set, the specified value is not valid for the specific attribute that is being set.
0x000002c2	706	GSK_CONFLICTING_VALIDATION_SETTING	Conflicting parameters were set for additional certificate validation.
0x000002c3	707	GSK_AES_UNSUPPORTED	The AES cryptographic algorithm is not supported.
0x000002c4	708	GSK_PEERID_LENGTH_ERROR	The PEERID does not have the correct length.
0x000002c5	709	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_OFF	The particular cipher is not allowed when FIPS mode of operation is off.
0x000002c6	710	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_ON	No approved FIPS ciphers are selected in FIPS mode of operation.
0x00000641	1601	GSK_TRACE_STARTED	The trace started successfully.
0x00000642	1602	GSK_TRACE_STOPPED	The trace stopped successfully.
0x00000643	1603	GSK_TRACE_NOT_STARTED	No trace file was previously started so it cannot be stopped.
0x00000644	1604	GSK_TRACE_ALREADY_STARTED	Trace file is started so it cannot be started again.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid, full-path file name.

Glosariusz

Ten glosariusz zawiera terminy i definicje dla produktu IBM Spectrum Protect oraz powiązanych produktów.

W tym glosariuszu używane są następujące odsyłacze:

- *Patrz* odsyła z niepreferowanego do preferowanego terminu lub ze skrótu do pełnej formy terminu.
- *Patrz także* odsyła do terminu pokrewnego lub przeciwnego.

Inne terminy i definicje są dostępne w serwisie IBM Terminology.

ABCDEFGHIJKLMNPRŚTUVWZ

A

ACK

Patrz potwierdzenie.

ACL

Patrz lista kontroli dostępu.

adaptacyjne tworzenie kopii zapasowej fragmentu pliku

Typ operacji tworzenia kopii zapasowej, w którym do serwera przesyłane są tylko te fragmenty pliku, które uległy zmianie; nie jest przesyłany cały plik. Adaptacyjne tworzenie kopii zapasowej fragmentu pliku zmniejsza ruch w sieci i zwiększa szybkość operacji tworzenia kopii zapasowej.

administrator

Osoba odpowiedzialna za zadania administracyjne, takie jak autoryzacja dostępu i zarządzanie treścią. Administratorzy mogą także nadawać użytkownikom uprawnienia różnego poziomu.

adres IP

Unikalny adres urządzenia lub jednostki logicznej w sieci zgodny ze standardem IP (Internet Protocol).

agent LVSA (Logical Volume Snapshot Agent)
Oprogramowanie pełniące rolę dostawcy obrazu stanu w celu utworzenia obrazu stanu woluminu logicznego podczas tworzenia kopii zapasowej obrazu stanu podczas pracy.

agent pamięci masowej
Program umożliwiający bezpośrednie tworzenie kopii zapasowej danych klienta i ich odtwarzanie w pamięci masowej podłączonej do sieci SAN.

agregat
Obiekt zapisany w jednej lub większej liczbie pul pamięci masowej składający się z grupy plików logicznych wspólnie pogrupowanych. Patrz także plik logiczny, plik fizyczny.

akceptor klienta
Usługa udostępniająca przeglądarkom WWW aplet Java klienta WWW. W systemach Windows akceptor klienta jest instalowany i uruchamiany jako usługa. W systemach AIX, UNIX i Linux akceptor klienta jest uruchamiany jako demon.

aktywny system plików
System plików, do którego dodano zarządzanie pamięcią masową. Gdy obsługiwane jest zarządzanie pamięcią masową, zadania dla aktywnego systemu plików obejmują: migrację automatyczną, uzgadnianie, migrację wybiórczą i przywracanie. Patrz także nieaktywny system plików.

aktywny zestaw strategii
Aktywowany zestaw strategii zawierający reguły strategii obecnie używane przez wszystkie węzły klienckie przypisane do domeny strategii. Patrz także domena strategii, zestaw strategii.

aktywować
Sprawdzić zawartość zestawu strategii, a następnie ustawić ten zestaw strategii jako aktywny.

algorytm Nagle'a
Algorytm zmniejszający przeciążenie w sieciach TCP/IP przez łączenie mniejszych pakietów i wysyłanie ich razem.

aplikacyjny interfejs programistyczny zarządzania pamięcią masową danych (DSMAPI)
Zestaw funkcji i semantyka służąca do monitorowania zdarzeń dotyczących plików, zarządzania danymi w plikach i konserwacji tych danych. W środowisku programu interfejs DSMAPI używa zdarzeń do powiadamiania aplikacji zarządzających danymi o operacjach wykonywanych na plikach, przechowuje dowolne atrybuty dotyczące plików, obsługuje obszary zarządzane w plikach, a także steruje dostępem do obiektu pliku za pomocą uprawnień dostępu DSMAPI.

archiwizować
Kopiować programy, dane lub pliki na inny nośnik pamięci masowej zwykle w celu długoterminowego przechowywania lub ze względów bezpieczeństwa. Patrz także pobieranie.

atrybut rozszerzony (extended attribute - EA)
Nazwy lub pary wartości powiązane z plikami lub katalogami. Istnieją trzy klasy atrybutów rozszerzonych: atrybuty użytkownika, atrybuty systemowe i atrybuty zaufane.

AutoFS
Patrz system plików podłączany automatycznie.

automatyczne wykrywanie
Funkcja wykrywająca, raportująca i aktualizująca numer seryjny dysku lub biblioteki w bazie danych, gdy definiowana jest ścieżka z serwera lokalnego.

autoryzowany użytkownik
Użytkownik posiadający uprawnienia administratora klienta na stacji roboczej. Ten użytkownik zmienia hasła, dokonuje rejestracji otwartych i usuwa obszary plików.

B

baza danych plików wstępnie zmigrowanych

Baza danych zawierająca informacje o wszystkich plikach wstępnie zmigrowanych do pamięci masowej serwera.

biblioteka

1. Repozytorium dla nośników odłączalnych, takich jak dyski magnetyczne i taśmy magnetyczne.
2. Kolekcja zawierająca co najmniej jeden napęd i być może urządzenia automatyzujące (w zależności od typu biblioteki), która umożliwia dostęp do woluminów pamięci masowej.

biblioteka taśm

Zestaw urządzeń i narzędzi do obsługi środowiska taśm dla danej instalacji. Biblioteka taśm może zawierać stelaże przemysłowe dla taśmowej pamięci masowej, mechanizmy do automatycznego podłączania taśm, zestaw napędów taśm i zestaw powiązanych woluminów taśm podłączonych do tych napędów.

biblioteka współużytkowana

Urządzenie biblioteki używane przez wiele serwerów menedżerów pamięci masowej.

biblioteka zewnętrzna

Kolekcja napędów, które są zarządzane przez system zarządzania nośnikami inny niż serwer zarządzania pamięcią masową.
bieżąca kopia zapasowa
Pełna kopia zapasowa, w której pliki dzienników transakcji nie są usuwane, aby procedury tworzenia kopii zapasowych używające przyrostowych lub różnicowych kopii zapasowych nie zostały zakłócone.
buforować
Umieszczać zduplikowaną kopię pliku na nośniku o dostępie bezpośrednim, gdy serwer migruje ten plik do innej puli pamięci masowej w hierarchii.

C

CAD

Patrz demon akceptora klienta.

centralny program planujący

Funkcja umożliwiająca administratorowi zaplanowanie operacji klienta i komend administracyjnych. Operacje można zaplanować w taki sposób, aby były wykonywane okresowo lub w określonej dacie. Patrz także harmonogram komendy administracyjnej, harmonogram klienta.

centrum przetwarzania danych

W środowisku zwirtualizowanym jest to kontener, który przechowuje hosty, klastry, sieci i składnice danych.

czas dostępu do pliku

W systemach AIX, UNIX lub Linux: czas uzyskania ostatniego dostępu do pliku.

czas przechowywania

Mierzony w dniach czas, przez jaki kopie zapasowe i archiwalne plików są przechowywane w puli pamięci masowej, zanim zostaną usunięte. Czas przechowywania jest wyznaczany przez atrybuty grupy kopii i domyślne okresy karencji dla domeny.

częstotliwość

Atrybut grupy kopii określający minimalny odstęp czasu między tworzeniem przyrostowych kopii zapasowych (w dniach).

D

dane wiodące

Bajty danych z początkowej części zmigrowanego pliku, które są zapisane w pliku pośredniczącym odpowiadającym danemu plikowi w lokalnym systemie plików. Ilość danych wiodących zapisanych w pliku pośredniczącym zależy od określonej wielkości pliku pośredniczącego.

deduplikacja

Patrz także deduplikacja danych.

deduplikacja danych

Metoda zmniejszania ilości potrzebnej pamięci masowej przez eliminowanie nadmiarowych danych. Na nośniku pamięci masowej przechowywana jest tylko jedna instancja danych. Pozostałe instancje tych samych danych są zastępowane wskaźnikami do przechowywanej instancji. Patrz także deduplikacja podczas wstawiania, deduplikacja danych z przetwarzaniem końcowym.

deduplikacja danych z przetwarzaniem końcowym

Metoda zmniejszania ilości potrzebnej pamięci masowej przez eliminowanie nadmiarowych danych. Dane są najpierw zapisywane w puli pamięci masowej, następnie są identyfikowane zduplikowane dane, na koniec odzyskiwany jest wolny obszar w puli pamięci masowej. Patrz także deduplikacja danych, deduplikacja podczas wstawiania.

deduplikacja podczas wstawiania

Metoda zmniejszania ilości potrzebnej pamięci masowej przez eliminowanie nadmiarowych danych. Dane są deduplikowane w momencie zapisywania w kontenerze puli pamięci masowej. Patrz także deduplikacja danych, deduplikacja danych z przetwarzaniem końcowym.

demon

Program nienadzorowany, który w sposób nieprzerwany lub okresowy wykonuje funkcje, takie jak kontrola sieci.

demon akceptora klienta (CAD)

Patrz akceptor klienta.

demon kronikowania

W systemach AIX, UNIX i Linux: program śledzący zmiany w plikach znajdujących się w systemach plików.

demon monitorowania pamięci masowej

Demon, który sprawdza użycie miejsca we wszystkich systemach plików, dla których aktywna jest funkcja zarządzania pamięcią masową, i automatycznie uruchamia proces migracji progowej, jeśli stopień użycia miejsca w systemie plików osiąga lub przekracza próg górny.

domena

Grupa węzłów klienckich z co najmniej jednym zestawem strategii, która zarządza danymi lub zasobami pamięci masowej dla węzłów klienckich. Patrz także domena strategii.

domena klienta

Zestaw dysków, systemów plików lub woluminów wybranych przez użytkownika w celu utworzenia kopii zapasowej lub archiwalnej danych za pomocą klienta kopii zapasowych i archiwalnych.

domena strategii

Grupa użytkowników strategii z jednym lub większą liczbą zestawów strategii, która zarządza danymi lub zasobami pamięci masowej dla tych użytkowników. Użytkownikami są węzły klienckie powiązane z daną domeną strategii. Patrz także aktywny zestaw strategii, domena.

domyślna klasa zarządzania

Klasa zarządzania przypisana do zestawu strategii. Klasa ta jest wykorzystywana do zarządzania plikami kopii zapasowych lub plikami archiwum, gdy pliki te nie są jawnie powiązane z konkretną klasą zarządzania za pomocą listy włączyć/wykluczeń.

DRM

Patrz menedżer odtwarzania po awarii.

DSMAPI

Patrz aplikacyjny interfejs programistyczny zarządzania pamięcią masową danych.

dziennik aktywności

Dziennik, w którym rejestrowane są zwykłe komunikaty dotyczące aktywności generowane przez serwer. Komunikaty te zawierają informacje o operacjach serwera i klientów, takie jak czas uruchamiania sesji i błędy we/wy urządzeń.

dziennik błędów

Zestaw danych lub plik używany do rejestrowania informacji o błędach w produkcie lub systemie.

dziennik odtwarzania

Dziennik aktualizacji, które mają zostać zapisane w bazie danych. Za pomocą tego dziennika można przywrócić działanie systemu po awarii systemu lub nośników. Dziennik odtwarzania składa się z dziennika aktywnego (w tym kopii lustrzanej dziennika) i dzienników archiwalnych.

E

EA

Patrz atrybut rozszerzony.

EB

Patrz eksabajt.

EFS

Patrz szyfrowany system plików.

eksabajt (EB)

W kontekście procesora, rzeczywistej i wirtualnej pamięci masowej, a także ilości danych w kanale: dwa do potęgi 60 lub 1 152 921 504 606 846 976 bajtów. W kontekście dyskowej pamięci masowej i ilości przesyłanych danych: 1 000 000 000 000 000 000 bajtów.

F

fizyczny poziom zajętości

Ilość miejsca zajmowanego przez pliki fizyczne w puli pamięci masowej. Miejsce to uwzględnia miejsce niewykorzystane, które powstało po usunięciu plików logicznych z agregatów. Patrz także plik logiczny, logiczny poziom zajętości, plik fizyczny.

format rodzimy

Format danych zapisywanych w puli pamięci masowej bezpośrednio przez serwer. Patrz także nierodzimny format danych.

FSID

Patrz identyfikator obszaru plików.

FSM

Patrz migrator systemu plików.

G

GB

Patrz gigabajt.

General Parallel File System (GPFS)

Wysokowydajny system plików ze współużytkowanymi dyskami, który umożliwia dostęp do danych z węzłów w środowisku klastrowym. Patrz także zarządzanie cyklem życia informacji.

generowanie hasła

Proces tworzący i zapisujący nowe hasło w zaszyfrowanym pliku haseł, gdy stare hasło utraci ważność. Automatyczne generowanie hasła eliminuje konieczność wyświetlania monitu o wprowadzenie hasła.

gigabajt (GB)

(2) W kontekście pamięci procesora, rzeczywistej i wirtualnej pamięci masowej, a także ilości danych w kanale: dwa do potęgi 30 lub 1 073 741 824 bajty. W kontekście dyskowej pamięci masowej i ilości przesyłanych danych: 1 000 000 000 bajtów.

GPFS

Patrz General Parallel File System.

grupa kolokacji

Zdefiniowana przez użytkownika grupa węzłów klienckich, których dane są zapisane na minimalnej liczbie woluminów dzięki procesowi kolokacji.

grupa kopii

Obiekt strategii zawierający atrybuty sterujące sposobem generowania wersji kopii zapasowych lub kopii archiwalnych, miejscem, w którym wersje kopii zapasowych lub kopie archiwalne są początkowo umieszczane, a także momentem utraty ważności przez wersje kopii zapasowych lub kopie archiwalne. Grupa kopii należy do klasy zarządzania. Patrz także grupa kopii archiwalnych, grupa kopii zapasowych, wersja kopii zapasowej, klasa zarządzania.

grupa kopii archiwalnych

Obiekt strategii zawierający atrybuty sterujące generowaniem plików archiwum, ich miejscem docelowym i utratą ważności przez te pliki. Patrz także grupa kopii.

grupa kopii zapasowych

Obiekt strategii zawierający atrybuty sterujące generowaniem wersji kopii zapasowych plików, ich miejscem docelowym i utratą ważności przez te pliki. Grupa kopii zapasowych należy do klasy zarządzania. Patrz także grupa kopii.

GUID

Patrz unikalny identyfikator globalny.

H

harmonogram

Rekord bazy danych opisujący operacje klienta lub komendy administracyjne, które mają być przetworzone. Patrz także harmonogram komendy administracyjnej, harmonogram klienta.

harmonogram klienta

Rekord bazy danych opisujący planowane przetwarzanie operacji klienta w określonym czasie. Do operacji klienta należą: tworzenie kopii zapasowych i archiwalnych, odtwarzanie, pobieranie, wykonywanie komendy systemu operacyjnego klienta lub makra. Patrz także harmonogram komendy administracyjnej, centralny program planujący, harmonogram.

harmonogram komendy administracyjnej

Rekord bazy danych opisujący planowane przetwarzanie komendy administracyjnej w określonym czasie. Patrz także centralny program planujący, harmonogram klienta, harmonogram.

hierarchia pamięci masowej

Zdefiniowany przez administratora logiczny porządek podstawowych pul pamięci masowej. Porządek ten jest zwykle oparty na szybkości i pojemności urządzeń, których używają poszczególne pule pamięci. Hierarchia pamięci masowej jest zdefiniowana przez określenie następczej puli pamięci masowej w definicji puli pamięci masowej. Patrz także pula pamięci masowej.

hierarchiczne zarządzanie pamięcią masową (HSM)

Funkcja, która automatycznie dystrybuje dane i zarządza nimi na dyskach, taśmach lub dyskach i taśmach, traktując urządzenia o tych i potencjalnie innych typach jako poziomy w hierarchii pamięci masowej, które występują w segmentach od urządzeń szybkich i drogiej do urządzeń wolniejszych, tańszych i prawdopodobnie wymiennych. Głównym celem jest minimalizacja czasu dostępu do danych i maksymalizacja dostępnej pojemności. Patrz także klient hierarchicznego zarządzania pamięcią masową, przywracanie, hierarchia pamięci masowej.

HSM

Patrz hierarchiczne zarządzanie pamięcią masową.

I

i-węzeł

Wewnętrzna struktura opisująca pojedyncze pliki w systemach AIX, UNIX lub Linux. I-węzeł zawiera węzeł, typ, właściciela i położenie pliku.

identyfikator obszaru plików (file space ID - FSID)

Unikalny identyfikator liczbowy, jaki serwer przypisuje do obszaru plików, gdy jest on zapisywany w pamięci masowej serwera.

ILM

Patrz zarządzanie cyklem życia informacji.

K

kB

Patrz kilobajt.

kilobajt (kB)

W kontekście pamięci procesora, rzeczywistej i wirtualnej pamięci masowej, a także ilości danych w kanale: dwa do potęgi 10 lub 1 024 bajty. W kontekście dyskowej pamięci masowej i ilości przesyłanych danych: 1 000 bajtów.

klasa uprawnień

Poziom uprawnień nadanych administratorowi. Klasa uprawnień określa czynności administracyjne, które administrator może wykonywać. Patrz także uprawnienie, klasa uprawnień węzła, klasa uprawnień operatora, klasa uprawnień strategii, klasa uprawnień pamięci masowej, klasa uprawnień systemowych.

klasa uprawnień administracyjnych

Patrz klasa uprawnień.

klasa uprawnień operatora

Klasa uprawnień pozwalająca administratorowi na wyłączenie lub zatrzymanie serwera, włączenie serwera, anulowanie procesów serwera i zarządzanie nośnikami wymiennymi. Patrz także klasa uprawnień.

klasa uprawnień pamięci masowej

Klasa uprawnień dająca administratorowi uprawnienia do sterowania sposobem przydzielania i używania zasobów pamięci serwera, np. monitorowania bazy danych, dziennika aktywności i pamięci masowej serwera. Patrz także klasa uprawnień.

klasa uprawnień strategii

Klasa uprawnień zezwalająca administratorowi na zarządzanie obiektami strategii, rejestrowanie węzłów klienckich i planowanie operacji klientów dla węzłów klienckich. Uprawnienia te mogą być ograniczone do określonych domen strategii. Patrz także klasa uprawnień.

klasa uprawnień systemowych

Klasa uprawnień dająca administratorowi uprawnienia do wykonywania wszystkich komend serwera. Patrz także klasa uprawnień.

klasa uprawnień węzła

Klasa uprawnień dająca administratorowi uprawnienia do zdalnego dostępu do klientów kopii zapasowych i archiwalnych dla określonego węzła klienckiego lub dla wszystkich klientów w domenie strategii. Patrz także klasa uprawnień.

klasa urządzeń

Nazwany zestaw parametrów stosowanych do grupy urządzeń pamięci masowej. Każda klasa urządzeń ma unikalną nazwę i reprezentuje urządzenie takie jak dysk, plik, dysk optyczny lub taśma.

klasa zarządzania

Obiekt strategii, który użytkownicy mogą powiązać z każdym plikiem w celu określenia sposobu używanego przez serwer do zarządzania plikiem. Klasa zarządzania może zawierać grupę kopii zapasowych, grupę kopii archiwalnych i atrybuty zarządzania pamięcią masową. Patrz także powiązać, grupa kopii, klient hierarchicznego zarządzania pamięcią masową, zestaw strategii, ponowne wiązanie.

klient

Program lub komputer żądający usług z serwera. Patrz także serwer.

klient administracyjny

Program działający na serwerze plików, stacji roboczej lub komputerze mainframe używany przez administratorów do sterowania serwerem i jego monitorowania. Patrz także klient kopii zapasowych i archiwalnych.

klient aplikacji

Program zainstalowany w systemie w celu zabezpieczenia aplikacji. Serwer udostępnia usługi tworzenia kopii zapasowej klientowi aplikacji.

klient biblioteki

Serwer używający komunikacji serwer-serwer do uzyskania dostępu do biblioteki zarządzanej przez inny serwer zarządzania pamięci masowej. Patrz także menedżer biblioteki.

klient hierarchicznego zarządzania pamięcią masową (klient HSM)

Program kliencki działający razem z serwerem w celu udostępnienia systemowi hierarchicznego zarządzania pamięcią masową (HSM). Patrz także hierarchiczne zarządzanie pamięcią masową, klasa zarządzania.

klient HSM

Patrz klient hierarchicznego zarządzania pamięcią masową.

klient kopii zapasowych i archiwalnych

Program działający na stacji roboczej lub na serwerze plików udostępniający użytkownikom narzędzia do tworzenia kopii zapasowych plików, archiwizowania i odtwarzania plików, a także ich pobierania. Patrz także klient administracyjny.

klient/serwer

Termin odnoszący się do modelu interakcji w przetwarzaniu danych rozproszonych, w którym program znajdujący się na jednym komputerze wysyła żądania do programu znajdującego się na innym komputerze i czeka na odpowiedź. Program żądający jest zwany klientem, a program odpowiadający jest zwany serwerem.

kod pośredniczący

Skrót w systemie plików Windows generowany przez klienta HSM dla zmigrowanego pliku umożliwiający użytkownikom przezroczysty dostęp do plików. Kod pośredniczący jest reprezentacją zmigrowanego pliku w postaci pliku rzadkiego z przyłączonym punktem ponownej analizy.

kolekcja zestawów kopii zapasowych

Grupa zestawów kopii zapasowych tworzonych w tym samym czasie i mających te same nazwy zestawów kopii zapasowych, nazwy woluminów, opisy i klasy urządzeń. Serwer identyfikuje każdy zestaw kopii zapasowych za pomocą nazwy węzła, nazwy zestawu kopii zapasowych i typu pliku.

kolokacja

Proces utrzymywania wszystkich danych należących do obszaru plików jednego klienta, jednego węzła klienckiego lub grupy węzłów klienckich na minimalnej liczbie woluminów o dostępie sekwencyjnym w puli pamięci masowej. Kolokacja może zmniejszyć liczbę woluminów, do których trzeba uzyskać dostęp, gdy konieczne jest odtworzenie ogromnej ilości danych.

kompresja

Funkcja usuwająca powtarzające się znaki, spacje lub łańcuchy znaków albo dane binarne z przetwarzanych danych i zastępująca je znakami sterującymi. Kompresja zmniejsza ilość pamięci masowej wymaganej do przechowywania danych. Patrz także kompresja podczas wstawiania.

kompresja podczas wstawiania

Metoda redukowania zajmowanego miejsca. Powtarzające się znaki, spacje, łańcuchy znaków lub dane binarne są usuwane w momencie zapisywania danych do puli pamięci masowej kontenera. Patrz także kompresja.

konfiguracja korporacyjna

Metoda konfigurowania serwerów, aby administrator mógł dystrybuować konfigurację jednego z serwerów na inne serwery za pomocą komunikacji typu serwer-serwer. Patrz także menedżer konfiguracji, serwer zarządzany, profil, subskrypcja.

kontener

Położenie pamięci masowej danych, na przykład plik, katalog lub urządzenie. Patrz także pula pamięci masowej kontenera.

kontrolować

Sprawdzać logiczne niespójności między informacjami posiadanymi przez serwer a rzeczywistym stanem systemu. Menedżer pamięci masowej może kontrolować informacje dotyczące takich elementów jak woluminy, biblioteki i licencje. Na przykład, gdy menedżer pamięci masowej kontroluje wolumin, serwer wyszukuje niespójności między informacjami o plikach, dla których utworzono kopie zapasowe lub archiwalne, zapisanymi w bazie danych a rzeczywistymi danymi powiązаныmi z każdą kopią zapasową lub archiwalną w pamięci masowej serwera.

konwersacja

Połączenie między dwoma programami zrealizowane za pomocą sesji, które umożliwia im wzajemną komunikację podczas przetwarzania transakcji.

kopia archiwalna

Plik lub grupa plików zarchiwizowanych w pamięci masowej serwera.

kopia rozmyta

Wersja kopii zapasowej lub kopii archiwalnej pliku, która może nie odzwierciedlać dokładnie oryginalnej zawartości tego pliku, ponieważ kopia zapasowa lub archiwalna tego pliku została utworzona w czasie jego modyfikacji.

kopia zapasowa grupy

Tworzenie kopii zapasowej grupy zawierającej listę plików pochodzących z co najmniej jednego źródłowego obszaru plików.

kopia zapasowa obrazu stanu

Kopia zapasowa pełnego systemu plików lub surowego woluminu logicznego jako pojedynczego obiektu.

kopia zapasowa woluminu logicznego

Kopia zapasowa systemu plików lub woluminu logicznego jako pojedynczego obiektu.

kopiowanie w tle

Obraz stanu woluminu. Obraz stanu może być tworzony również wtedy, gdy aplikacje w systemie zapisują dane w woluminie.

L

LAN

Patrz sieć lokalna.

limit

1. Dla programu HSM w systemach AIX, UNIX i Linux: wyrażony w megabajtach limit ilości danych, które mogą zostać zmigrowane i zmigrowane wstępnie z systemu plików do pamięci masowej serwera.
2. Dla programu HSM w systemach Windows: zdefiniowany przez użytkownika limit miejsca zajmowanego przez przywrócone pliki.

limit czasu

Przydzielony odstęp czasu, w którym zdarzenie ma wystąpić lub zostać zakończone przed przerwaniem operacji.

limit połączeń

Maksymalna liczba woluminów, jakie mogą być równocześnie dostępne w danej klasie urządzeń. Limit połączeń określa maksymalną liczbę punktów połączenia. Patrz także punkt połączenia.

lista dołączeń i wykluczeń

Lista opcji, które włączają lub wykluczają wybrane pliki z operacji tworzenia kopii zapasowej. Opcja wykluczania (exclude) określa pliki, dla których kopia zapasowa nie będzie tworzona. Opcja włączania (include) określa pliki stanowiące wyjątek od reguł wykluczania lub przypisuje klasę zarządzania do pliku albo grupy plików dla usług tworzenia kopii zapasowych lub archiwalnych. Patrz także plik włączeń/wykluczeń.

lista kontroli dostępu (ACL)

W zagadnieniach związanych z bezpieczeństwem komputerów: lista powiązana z obiektem, która określa wszystkie podmioty mające dostęp do obiektu i prawa dostępu.

lista wykluczeń/włączeń

Patrz lista dołączeń i wykluczeń.

LOFS

Patrz wirtualny system plików LOFS.

logiczny poziom zajętości

Miejsce zajmowane przez pliki logiczne w puli pamięci masowej. Miejsce to nie uwzględnia wolnego miejsca powstałego wskutek usunięcia plików logicznych z plików agregatów i dlatego może być go mniej niż w przypadku fizycznego poziomu zajętości. Patrz także fizyczny poziom zajętości.

lokalne

1. Termin odnoszący się do urządzenia, pliku lub systemu, do którego dostęp jest uzyskiwany bezpośrednio z systemu użytkownika bez użycia linii komunikacyjnej.
2. W produktach HSM: termin odnoszący się do miejsca docelowego migrowanych plików, które są przenoszone. Patrz także zdalne.

lokalny wolumin w tle

Dane zapisane w woluminach w tle znajdujących się w dyskowym podsystemie pamięci masowej.

LUN

Patrz numer jednostki logicznej.

LVSA

Patrz agent LVSA.

M

macierzysty system plików

System plików, z którego dany plik został zmigrowany. Kiedy plik jest przywracany, jest on umieszczany z powrotem w macierzystym systemie plików.

maksymalna jednostka transmisji (MTU)

Największy możliwy blok danych, który może być wysłany za pomocą danego nośnika fizycznego w jednej ramce. Na przykład maksymalną jednostką transmisji dla sieci Ethernet jest 1500 bajtów.

MB

Patrz megabajt.

megabajt (MB)

W kontekście pamięci procesora, rzeczywistej i wirtualnej pamięci masowej, a także ilości danych w kanale: dwa do potęgi 20 lub 1 048 576 bajtów. W kontekście dyskowej pamięci masowej i ilości przesyłanych danych: 1 000 000 bajtów.

menedżer biblioteki

Serwer sterujący operacjami urządzenia, gdy wiele serwerów zarządzania pamięcią masową współużytkuje urządzenie pamięci masowej. Patrz także klient biblioteki.

menedżer konfiguracji

Serwer dystrybuujący informacje konfiguracyjne, takie jak strategie i harmonogramy, do serwerów zarządzanych zgodnie z ich profilami. Informacje konfiguracyjne mogą obejmować strategie i harmonogramy. Patrz także konfiguracja korporacyjna, serwer zarządzany, profil.

menedżer odtwarzania po awarii (DRM)

Funkcja pomagająca w przygotowaniu i użyciu planu odtwarzania po awarii dla serwera.

metadane

Dane opisujące charakterystykę danych; dane opisowe.

metoda komunikacji

Metoda, za pomocą której klient i serwer wymieniają informacje. Patrz także protokół TCP/IP (Transmission Control Protocol/Internet Protocol).

miejsce docelowe

Atrybut grupy kopii lub klasy zarządzania określający podstawową pulę pamięci masowej, w której zostanie utworzona kopia zapasowa lub archiwalna pliku klienta lub do której zostanie on zmigrowany. Patrz także pula pamięci masowej kopii.

migracja

Proces przenoszenia danych z jednego miejsca w inne lub aplikacji z jednego systemu komputerowego do drugiego.

migracja

Proces przenoszenia danych lub aplikacji z jednego systemu komputerowego do drugiego.

migracja automatyczna
Proces używany do automatycznego przenoszenia plików z lokalnego systemu plików lub pamięci podręcznej z uwzględnieniem opcji i ustawień wybranych przez administratora na stacji roboczej. Patrz także migracja na żądanie, migracja progowa.

migracja na żądanie
Proces używany do reagowania na sytuację braku miejsca w systemie plików, dla którego aktywne jest hierarchiczne zarządzanie pamięcią masową (HSM). Pliki są migrowane do serwera pamięci masowej do momentu, gdy użycie miejsca spadnie do dolnego progu ustawionego dla systemu plików. Jeśli progi górny i dolny są takie same, migrowany jest jeden plik. Patrz także migracja automatyczna, migracja wybiórcza, migracja progowa.

migracja progowa
Proces przenoszenia plików z lokalnego systemu plików do pamięci masowej serwera z uwzględnieniem wartości progów górnego i dolnego zdefiniowanych dla systemu plików. Patrz także migracja automatyczna, migracja na żądanie, zadanie migracji, migracja wybiórcza.

migracja wstępna
Proces kopiowania plików zakwalifikowanych do migracji do pamięci masowej serwera, w którym oryginalne pliki są pozostawiane w lokalnym systemie plików w stanie niezmienionym.

migracja wybiórcza
Proces kopiowania plików wybranych przez użytkownika z lokalnego systemu plików do pamięci masowej serwera i zastępowania ich plikami pośredniczącymi w lokalnym systemie plików. Patrz także migracja na żądanie, migracja progowa.

migrator systemu plików (file system migrator - FSM)
Rozszerzenie jądra przechwytyjące wszystkie operacje systemu plików i udostępniające wszelką wymaganą obsługę zarządzania pamięcią masową. Jeśli nie jest wymagana żadna obsługa zarządzania pamięcią masową, operacja jest przekazywana do systemu operacyjnego, który wykonuje zwykłe funkcje. Migrator systemu plików jest podłączany przez system plików, gdy do tego systemu plików dodawane jest zarządzanie pamięcią masową.

migrowany plik
Plik, który został skopiowany z lokalnego systemu plików do pamięci masowej. W przypadku klientów HSM w systemach UNIX lub Linux ten plik jest zastępowany plikiem pośredniczącym w lokalnym systemie plików. W systemach Windows utworzenie pliku pośredniczącego jest opcjonalne. Patrz także stan pliku, plik wstępnie zmigrowany, plik rezydentny, plik pośredniczący.

moment zatwierdzenia
Zadany moment, gdy dane uważa się za spójne.

MTU
Patrz maksymalna jednostka transmisji.

N

narzędzie przenoszenia danych
Narzędzie przenoszące dane w imieniu serwera. Narzędziem przenoszenia danych jest serwer plików pamięci masowej NAS.

natychmiastowe odtwarzanie VSS
Operacja, która odtwarza dane z lokalnego obrazu stanu. Obraz stanu jest kopią zapasową VSS, która znajduje się w lokalnym woluminie w tle. Operacja odtwarzania pobiera dane za pomocą metody odtwarzania wspieranej sprzętowo (na przykład FlashCopy).

nazwa węzła
Unikalna nazwa identyfikująca stację roboczą, serwer plików lub komputer PC na serwerze.

nazwa WWN
64-bitowy unikalny identyfikator nazwy (bez znaku).

NDMP
Patrz Network Data Management Protocol.

NetBIOS (Network Basic Input/Output System)
Standardowy interfejs sieci i komputerów osobistych używany w sieciach lokalnych do udostępniania funkcji obsługi komunikatów, serwera wydruków i serwera plików. Aplikacje używające interfejsu NetBIOS nie muszą obsługiwać szczegółów protokołów sterowania łączem danych (DLC) w sieci LAN.

Network Basic Input/Output System
Patrz NetBIOS.

Network Data Management Protocol (NDMP)
Protokół umożliwiający aplikacjom zarządzania pamięcią masową sterowanie operacjami tworzenia i odtwarzania kopii zapasowych serwera plików obsługującego protokół NDMP bez instalowania oprogramowania uzyskanego u dostawcy na tym serwerze plików.

nieaktywny system plików
System plików, w którym dezaktywowano zarządzanie pamięcią masową. Patrz także aktywny system plików.

nieobciążające żądanie kopii zapasowej VSS

Operacja tworzenia kopii zapasowej używająca dostawcy sprzętu Usługi kopiowania woluminów w tle (VSS) firmy Microsoft (zainstalowanego w systemie alternatywnym) do przeniesienia danych na serwer. Ten typ operacji tworzenia kopii zapasowej przenosi obciążenie związane z operacją tworzenia kopii zapasowej z systemu produkcyjnego do innego systemu.

nierodzimymy format danych

Format danych zapisywanych w puli pamięci masowej inny niż format, którego serwer używa podczas wykonywania operacji. Patrz także format rodzimy.

numer i-węzła

Numer określający konkretny plik i-węzła w systemie plików.

numer jednostki logicznej (LUN)

W standardzie Small Computer System Interface (SCSI): unikalny identyfikator używany do odróżniania urządzeń, z których każde jest jednostką logiczną (LU).

O

obiekt reliktowy

Niewielki podzbiór atrybutów usuniętego obiektu. Obiekt reliktowy jest przechowywany przez określony czas i na koniec podanego okresu obiekt ten jest trwale usuwany.

obiekt zarządzany

Definicja w bazie danych serwera zarządzanego, która została rozdystrybuowana na ten serwer zarządzany przez menedżera konfiguracji. Gdy serwer zarządzany zasubskrybuje profil, wszystkie obiekty powiązane z tym profilem staną się obiektami zarządzanymi w bazie danych serwera zarządzanego.

obraz

System plików lub surowy wolumin logiczny traktowany jako jeden obiekt, dla którego tworzona jest kopia zapasowa.

obraz stanu

Kopia zapasowa obrazu stanu zawierająca widok woluminu w zadanym momencie.

obraz stanu bazy danych

Pełna kopia zapasowa całej bazy danych na nośniku, który można przenieść fizycznie. Gdy tworzony jest obraz stanu bazy danych, nie jest przerywana bieżąca seria kopii zapasowych bazy danych. Z obrazem stanu bazy danych nie mogą być powiązane przyrostowe kopie zapasowe bazy danych. Patrz także seria kopii zapasowych bazy danych, pełna kopia zapasowa.

obszar

Część pliku tworzona podczas procesu deduplikacji danych. W celu znalezienia duplikatów obszary są porównywane z obszarami innych plików w celu.

obszar plików

Logiczny obszar w pamięci masowej serwera zawierający grupę plików, dla których węzeł kliencki utworzył kopię zapasową lub archiwalną z jednej partycji logicznej, systemu plików lub wirtualnego punktu podłączenia. Węzły klienckie mogą odtwarzać, pobierać lub usuwać swoje obszary plików z pamięci masowej serwera. W pamięci masowej serwera pliki należące do jednego obszaru plików niekoniecznie są zapisane razem.

obszar plików z obsługą Unicode

Obszar plików o nazwie, która jest zgodna ze standardem Unicode kompatybilnym z dowolnymi ustawieniami narodowymi na stacjach roboczych, w wielu językach.

odbiornik

Repozytorium serwera zawierające dziennik komunikatów serwera i klienta w postaci zdarzeń. Odbiornikiem może być na przykład program użytkownika, konsola serwera lub dziennik aktywności. Patrz także zdarzenie.

odtworzenie

Kopiowanie informacji z kopii zapasowej do aktywnej pamięci masowej w celu ich użycia. Przykład: kopiowanie informacji z pamięci masowej serwera do klienckiej stacji roboczej.

odtworzenie pojedynczej skrzynki pocztowej

Patrz odtwarzanie skrzynki pocztowej.

odtworzenie skrzynki pocztowej

Funkcja odtwarzająca dane programu Microsoft Exchange Server (z kopii zapasowych programu IBM Data Protection for Microsoft Exchange) na poziomie skrzynki pocztowej lub elementów skrzynki pocztowej.

odtworzenie VSS

Funkcja używająca usługi VSS (Microsoft Volume Shadow Copy Service) do odtwarzania obrazów stanu zapisanych w pamięci masowej serwera. Obrazy te zostały wcześniej utworzone przez funkcję kopii zapasowej VSS i są odtwarzane w oryginalnej lokalizacji.

odzyskiwanie

Proces konsolidowania pozostałych danych z wielu woluminów o dostępie sekwencyjnym do mniejszej liczby nowych woluminów o dostępie sekwencyjnym.

okno uruchamiania

Przedział czasu, w którym musi być zainicjowany harmonogram.

okres oczekiwania na podłączenie
Maksymalna liczba minut, przez jaką serwer oczekuje na wykonanie żądania podłączenia woluminu wolumin nośników o dostępie sekwencyjnym, zanim anuluje to żądanie.

okres przechowywania kopii archiwalnej
Liczba dni, przez które menedżer pamięci masowej przechowuje zarchiwizowany plik, gdy serwer nie może ponownie powiązać tego pliku z odpowiednią klasą zarządzania. Patrz także powiązać.

okres przechowywania kopii zapasowych
Liczba dni, przez które menedżer pamięci masowej przechowuje kopię zapasową, jeśli serwer nie może ponownie powiązać tego pliku z odpowiednią klasą zarządzania.

okres przechowywania podłączenia
Maksymalna liczba minut, przez jaką serwer przechowuje podłączony wolumin nośników o dostępie sekwencyjnym, który nie jest używany, zanim go odłączy.

osierocony plik pośredniczący
Plik, dla którego nie znaleziono pliku zmigrowanego na serwerze udostępniającemu węzłowi klienckiemu usługi zarządzania pamięcią masową. Na przykład plik pośredniczący może być osierocony wtedy, gdy plik opcji systemowych klienta zostanie zmodyfikowany w taki sposób, że będzie nawiązywane połączenie z serwerem innym niż serwer, do którego ten plik został zmigrowany.

ośrodek chroniony
Patrz także ośrodek podstawowy.

ośrodek dodatkowy
Ośrodek fizyczny lub wirtualny, który składa się z zasobów sprzętowych, sieciowych i pamięci masowej, używany na wypadek awarii ośrodka podstawowego. W razie wystąpienia awarii w ośrodku podstawowym operacje mogą być kontynuowane w ośrodku dodatkowym. Patrz także ośrodek podstawowy.

ośrodek odtwarzania
Patrz ośrodek dodatkowy.

ośrodek podstawowy
Ośrodek fizyczny lub wirtualny, który składa się z zasobów sprzętowych, sieciowych i pamięci masowej. Zwykle produkcja odbywa się w ośrodku podstawowym. Dane mogą być replikowane do ośrodka dodatkowego na wypadek konieczności przetłoczenia awaryjnego lub odzyskiwania po awarii. Patrz także ośrodek dodatkowy.

P

pakiet
W zagadnieniach dotyczących przesyłania danych: sekwencja cyfr binarnych, w tym danych i sygnałów sterujących, przesyłanych i przetwarzanych jako całość.

pamięć masowa serwera
Pule pamięci podręcznej, podstawowe, kopii i danych aktywnych używane przez serwer do zapisywania plików użytkowników: wersji kopii zapasowych, kopii archiwalnych i plików zmigrowanych z węzłów klientów HSM (plików z zarządzaniem pamięcią). Patrz także pula danych aktywnych, pula pamięci masowej kontenera, pula pamięci masowej kopii, podstawowa pula pamięci masowej, wolumin puli pamięci masowej, wolumin.

partycja zarządzania obciążeniem (WPAR)
Partycja w pojedynczej instancji systemu operacyjnego.

pełna kopia zapasowa
Proces tworzenia kopii zapasowej całej bazy danych serwera. Pełna kopia zapasowa rozpoczyna nową serię kopii zapasowych bazy danych. Patrz także seria kopii zapasowych bazy danych, obraz stanu bazy danych, przyrostowa kopia zapasowa.

plan odtwarzania po awarii
Plik utworzony przez menedżera odtwarzania po awarii (DRM) zawierający informacje o sposobie odtwarzania systemów komputerowych po wystąpieniu awarii i skrypty, które można uruchomić w celu wykonania niektórych zadań odtwarzania. Plik ten zawiera informacje o oprogramowaniu i sprzęcie używanym przez serwer i miejscu, w którym znajduje się nośnik odtwarzania.

plik fizyczny
Plik zapisany w jednej lub większej liczbie pul pamięci masowej, składający się z jednego pliku logicznego lub grupy plików logicznych tworzących razem agregat. Patrz także agregat, plik logiczny, fizyczny poziom zajętości.

plik historii woluminów
Plik zawierający informacje o woluminach, które były używane przez serwer podczas operacji tworzenia kopii zapasowych baz danych i do eksportowania danych administratora, węzła, strategii lub serwera. W tym pliku ponadto znajdują się informacje o woluminach pul pamięci masowej o dostępie sekwencyjnym, które zostały dodane, ponownie wykorzystane lub usunięte. Informacje te są kopią informacji o woluminach zapisanych w bazie danych serwera.

plik konfiguracyjny urządzenia

1. W przypadku serwera: plik zawierający informacje o zdefiniowanych klasach urządzeń i - na niektórych serwerach - o zdefiniowanych bibliotekach i napędach. Informacje te to kopia informacji konfiguracyjnych urządzenia w bazie danych.
2. W przypadku agenta pamięci masowej: plik zawierający nazwę i hasło agenta pamięci masowej, informacje o serwerze zarządzającym bibliotekami i napędami podłączonymi do sieci SAN używanymi przez agenta pamięci masowej.

plik logiczny

Plik zapisany w jednej lub większej liczbie pul pamięci masowej serwera jako samodzielny plik lub jako część agregatu. Patrz także agregat, plik fizyczny, fizyczny poziom zajętości.

plik makra

Plik zawierający jedną lub więcej komend administracyjnych produktu IBM Spectrum Protect, który może być uruchomiony wyłącznie z klienta administracyjnego za pomocą komendy MACRO. Patrz także skrypt komend produktu IBM Spectrum Protect.

plik opcji

Plik zawierający opcje przetwarzania. Patrz także plik opcji systemowych klienta, plik opcji użytkownika klienta.

plik opcji klienta

Edytowalny plik określający serwer i metodę komunikacji oraz udostępniający konfigurację dla operacji tworzenia kopii zapasowych, archiwizowania, hierarchicznego zarządzania pamięcią masową i planowania.

plik opcji serwera

Plik zawierający ustawienia, które sterują różnymi operacjami serwera. Ustawienia te wpływają między innymi na procesy komunikacji, urządzenia i wydajność serwera.

plik opcji systemowych klienta

AIX, UNIX lub Linux zawierający zestaw opcji przetwarzania określających serwery, z którymi nawiązywane są połączenia w celu udostępnienia usług. Ponadto ten plik określa metody komunikacji i opcje dla operacji tworzenia kopii zapasowych, archiwizowania, hierarchicznego zarządzania pamięcią masową i planowania. Patrz także plik opcji użytkownika klienta, plik opcji.

plik opcji użytkownika klienta

Plik zawierający zestaw opcji przetwarzania używanych przez klientów w systemie. Ten zestaw może obejmować opcje określające serwer, z którym komunikuje się klient, a także opcje wpływające na operacje tworzenia kopii zapasowych, operacje archiwizowania, operacje hierarchicznego zarządzania pamięcią masową i operacje zaplanowane. Ten plik jest zwany również plikiem dsm.opt. Dla systemów AIX, UNIX lub Linux patrz także plik opcji systemowych klienta. Patrz także plik opcji systemowych klienta, plik opcji.

plik pamięci podręcznej

Obraz stanu woluminu logicznego utworzony przez agenta LVSA (Logical Volume Snapshot Agent). Bloki są zapisywane tuż przed modyfikacją podczas tworzenia kopii zapasowej obrazu stanu, a ich obszary logiczne są zapisywane w plikach pamięci podręcznej.

plik pośredniczący

Plik, który zastępuje oryginalny plik w lokalnym systemie plików, gdy zostanie on zmigrowany do pamięci masowej. Plik pośredniczący zawiera informacje niezbędne do przywrócenia zmigrowanego pliku z pamięci masowej serwera. Plik ten zawiera również dodatkowe informacje, których można użyć, aby nie trzeba było przywracać zmigrowanego pliku. Patrz także migrowany plik, plik rezydentny.

plik przedawniony

Zmigrowany lub zmigrowany wstępnie plik oznaczony jako plik, który ma utracić ważność i zostać usunięty z pamięci masowej. Jeśli plik pośredniczący lub oryginalna kopia wstępnie zmigrowanego pliku zostanie usunięta z systemu lokalnego albo jeśli oryginalna kopia wstępnie zmigrowanego pliku zostanie zaktualizowana, odpowiadający jej zmigrowany lub wstępnie zmigrowany plik zostanie oznaczony do utraty ważności i usunięcia przy następnym uruchomieniu uzgadniania.

plik rezydentny

W systemie Windows: kompletny plik w lokalnym systemie plików, którym może być także plik zmigrowany, ponieważ w pamięci masowej serwera istnieje zmigrowana kopia. W systemach UNIX lub Linux: kompletny plik w lokalnym systemie plików, który nie został zmigrowany lub zmigrowany wstępnie albo który został przywrócony z pamięci masowej serwera i zmodyfikowany.

plik rzadki

Plik większy niż ilość znajdujących się w nim danych, zawierający puste miejsca na dane, które zostaną dodane w przyszłości.

plik specjalny

W systemach AIX, UNIX i Linux: plik definiujący urządzenia dla systemu lub pliki specjalne tworzone przez procesy. Istnieją trzy podstawowe typy plików specjalnych: FIFO (first-in, first-out), blokowe i znakowe.

plik włączeń/wykluczeń

Plik zawierający instrukcje określające pliki, dla których ma być tworzona kopia zapasowa, a także powiązane klasy zarządzania, które mają być używane przez operacje tworzenia kopii zapasowych lub archiwalnych. Patrz także lista

dołączeń i wykluczeń.

plik wstępnie zmigrowany
Plik, który został skopiowany do pamięci masowej serwera, ale nie został zastąpiony plikiem pośredniczącym w lokalnym systemie plików. Identyczna kopia pliku istnieje zarówno w lokalnym systemie plików, jak i w pamięci masowej serwera. Pliki wstępnie zmigrowane występują w systemach plików UNIX i Linux, do których dodano zarządzanie pamięcią masową. Patrz także stan pliku, migrowany plik, plik rezydentny.

plik z zarządzaniem pamięcią
Plik migrowany z węzła klienckiego przez klienta HSM. Klient HSM na żądanie przywraca plik do węzła klienckiego.

plik zadania
Wygenerowany plik zawierający informacje konfiguracyjne dotyczące zadania migracji. Plik ten jest w formacie XML i może być tworzony, a także modyfikowany w kliencie HSM dla graficznego interfejsu użytkownika klienta Windows. Patrz także zadanie migracji.

plik zniszczony
Plik fizyczny, w którym wykryto błędy odczytu.

pobieranie
Kopiowanie zarchiwizowanych informacji z puli pamięci masowej do stacji roboczej w celu ich użycia. Operacja pobierania nie ma wpływu na wersję archiwalną zapisaną w puli pamięci masowej. Patrz także archiwizować.

podstawowa pula pamięci masowej
Nazwany zestaw woluminów lub kontenerów, których serwer używa do zapisywania wersji kopii zapasowych i archiwalnych plików, a także plików zmigrowanych z węzłów klienckich. Patrz także pula pamięci masowej kopii, pamięć masowa serwera, pula pamięci masowej, wolumin puli pamięci masowej.

ponowne wiązanie
Kojarzenie wszystkich wersji kopii zapasowych pliku z nową nazwą klasy zarządzania. Na przykład plik, który ma aktywną wersję kopii zapasowej, jest powiązany ponownie, gdy dla nowszej wersji tego pliku zostanie utworzona kopia zapasowa z użyciem innego powiązania klasy zarządzania. Patrz także powiązać, klasa zarządzania.

porcja
Kontener pamięci masowej w chmurze używany przez usługę Amazon Simple Storage Service (Amazon S3).

potok nazwany
Typ komunikacji międzyprocesowej pozwalającej na przesyłanie strumieni danych między równorzędnymi procesami węzłów sieci, takich jak klient i serwer.

potwierdzenie (ACK)
Przesłanie znaków potwierdzenia odbioru jako pozytywnej odpowiedzi na przesłanie danych.

powiązać
Skojarzyć plik z nazwą klasy zarządzania. Patrz także okres przechowywania kopii archiwalnej, klasa zarządzania, ponowne wiązanie.

powiązania profilu
W menedżerze konfiguracji: relacja zdefiniowana między profilem a obiektem, takim jak domena strategii. Powiązania profilu definiują informacje konfiguracyjne dystrybuowane do serwera zarządzanego, gdy profil ten zostanie zasubskrybowany przez serwer.

powiązanie
Zdefiniowana relacja między węzłem klienckim a harmonogramem klienta. Powiązanie określa nazwę harmonogramu, nazwę domeny strategii, do której należy ten harmonogram, a także nazwę węzła klienckiego wykonującego zaplanowane operacje.

profil
Nazwana grupa informacji konfiguracyjnych, która może być dystrybuowana z menedżera konfiguracji po zasubskrybowaniu przez serwer zarządzany. Informacje konfiguracyjne mogą obejmować: zarejestrowane identyfikatory administratorów, strategię, harmonogramy klientów, zestawy opcji klientów, harmonogramy administrowania, skrypty komend menedżera pamięci masowej, definicje serwerów i grup serwerów. Patrz także menedżer konfiguracji, konfiguracja korporacyjna, serwer zarządzany.

protokół komunikacyjny
Zestaw zdefiniowanych interfejsów umożliwiających komputerom wzajemną komunikację.

protokół TCP/IP (Transmission Control Protocol/Internet Protocol)
Zestaw protokołów komunikacyjnych, stanowiących standard branżowy i nie będących protokołami prawnie zastrzeżonymi, zapewniających niezawodne połączenia na całej trasie między aplikacjami za pomocą sieci sprzęgających o różnych typach. Patrz także metoda komunikacji.

próg migracji
Dolne i górne pojemności puli pamięci masowej lub systemów plików wyrażone jako wartości procentowe, przy których migracja ma być uruchamiana i zatrzymywana.

próg odzyskiwania
Procentowa wartość pamięci masowej, jaka musi istnieć na woluminie nośników o dostępie sekwencyjnym, aby serwer mógł odzyskać ten wolumin. Pamięć staje się odzyskiwalna w miarę utraty ważności przez pliki lub usuwania plików.

przedrostek woluminu taśm

Kwalifikator wysokiego poziomu nazwy pliku lub nazwy zestawu danych na standardowej etykiecie taśmy.

przenoszenie danych bez obciążania sieci LAN
Przenoszenie danych klienta między systemem klienta i urządzeniem pamięci masowej w sieci SAN z pominięciem sieci lokalnej.

przepustowość
W zagadnieniach związanych z zarządzaniem pamięcią masową: łączna liczba bajtów związanych z obciążeniem, z wyłączeniem narzutu, dla których wykonywana jest operacja tworzenia lub odtwarzania kopii zapasowej, podzielona przez czas trwania tej operacji.

przesyłanie danych bez obciążania sieci LAN
Patrz przenoszenie danych bez obciążania sieci LAN.

przyrostowa kopia zapasowa
Proces tworzenia kopii zapasowej plików i katalogów lub kopiowania stron w bazie danych, które są nowe lub które uległy zmianie od chwili utworzenia ostatniej pełnej lub przyrostowej kopii zapasowej. Patrz także wybiórcze tworzenie kopii zapasowych.

przywracanie
Kopiowanie zmigrowanego pliku z pamięci masowej serwera z powrotem do macierzystego systemu plików za pomocą klienta zarządzania pamięcią masową. Patrz także przywracanie wybiórcze.

przywracanie przezroczyste
Proces używany w celu automatycznego przywrócenia zmigrowanego wcześniej pliku do stacji roboczej lub serwera plików w chwili uzyskania dostępu do pliku. Patrz także przywracanie wybiórcze.

przywracanie wybiórcze
Proces kopiowania plików wybranych przez użytkownika z pamięci masowej serwera do lokalnego systemu plików. Patrz także przywracanie, przywracanie przezroczyste.

pula danych aktywnych
Nazwany zestaw woluminów puli pamięci masowej zawierający tylko aktywne wersje danych kopii zapasowej klienta. Patrz także pamięć masowa serwera, pula pamięci masowej, wolumin puli pamięci masowej.

pula pamięci masowej
Zestaw woluminów pamięci masowej lub kontenerów, który jest miejscem docelowym używanym do zapisywania danych klientów. Patrz także pula danych aktywnych, pula pamięci masowej kontenera-chmury, pula pamięci masowej kopii, pula pamięci masowej kontenera-katalogu, podstawowa pula pamięci masowej, hierarchia pamięci masowej.

pula pamięci masowej kontenera
Podstawowa pula pamięci masowej, której serwer używa do przechowywania danych. Dane są przechowywane w kontenerach w katalogach systemu plików lub w pamięci masowej w chmurze. Dane są deduplikowane, jeśli to konieczne, ponieważ serwer zapisuje je w puli pamięci masowej. Patrz także pula pamięci masowej kontenera-chmury, kontener, pula pamięci masowej kontenera-katalogu.

pula pamięci masowej kontenera-chmury
Pula pamięci masowej, której serwer używa do przechowywania danych w pamięci masowej w chmurze. Pamięć masowa w chmurze może znajdować się w siedzibie lub poza siedzibą przedsiębiorstwa. Patrz także pamięć masowa kontenera, pula pamięci masowej katalogu-kontenera, pula pamięci masowej.

pula pamięci masowej kontenera-katalogu
Pula pamięci masowej, której serwer używa do przechowywania danych w katalogach pamięci masowej. Dane, które są przechowywane w puli pamięci masowej kontenera-katalogu, mogą podlegać deduplikacji podczas wstawiania lub po stronie klienta. Patrz także pula pamięci masowej kontenera-chmury, pula pamięci masowej kontenera, pula pamięci masowej kontenera-kopii, pula pamięci masowej.

pula pamięci masowej kontenera-kopii
Pula pamięci masowej, której serwer używa do przechowywania kopii obszarów z puli pamięci masowej w chmurze katalogu-kontenera. Kopie są używane do naprawy uszkodzeń w puli pamięci masowej kontenera-katalogu. Pule pamięci masowej kontenera-kopii używają nośników sekwencyjnych, takich jak taśma. Patrz także pula pamięci katalogu-kontenera.

pula pamięci masowej kopii
Nazwany zestaw woluminów zawierających kopie plików znajdujących się w podstawowych pulach pamięci masowej. Pule pamięci masowej kopii są używane tylko do tworzenia kopii zapasowych danych zapisanych w podstawowych pulach pamięci masowej. Pula pamięci masowej kopii nie może być miejscem docelowym dla grupy kopii zapasowych, grupy kopii archiwalnych ani klasy zarządzania (dla plików z zarządzaniem pamięcią). Patrz także miejsce docelowe, podstawowa pula pamięci masowej, pamięć masowa serwera, pula pamięci masowej, wolumin puli pamięci masowej.

punkt podłączenia
Napęd logiczny, za pomocą którego uzyskuje się dostęp do woluminów należących do klasy urządzeń o dostępie sekwencyjnym. Dla nośników wymiennych, takich jak taśmy, punkt podłączenia jest napędem logicznym powiązany z napędem fizycznym. Dla urządzeń o typie plikowym, punkt podłączenia jest napędem logicznym powiązany ze strumieniem we/wy. Patrz także limit podłączeń.

randomizacja
Proces dystrybuowania czasów uruchamiania harmonogramów dla różnych klientów w obrębie określonej wartości procentowej okna uruchamiania harmonogramu.

reguła autoryzacji
Specyfikacja zezwalająca innemu użytkownikowi na odtwarzanie lub pobieranie plików użytkownika z pamięci masowej.

reguła uwierzytelniania
Specyfikacja określająca, że inny użytkownik może odtwarzać lub pobierać pliki z pamięci masowej.

rejestr
Repozytorium zawierające informacje o dostępie i konfiguracji dotyczące użytkowników, systemów i oprogramowania.

rejestracja otwarta
Proces rejestrowania, w którym użytkownicy mogą sami rejestrować na serwerze swoje stacje robocze jako węzły klienckie. Patrz także rejestracja zamknięta.

rejestracja zamknięta
Proces rejestracji, w którym tylko administrator może rejestrować stacje robocze jako węzły klienckie na serwerze. Patrz także rejestracja otwarta.

rejestrowanie
Definiowanie węzła klienckiego lub identyfikatora administratora, który może uzyskiwać dostęp do serwera.

rejestrowanie korporacyjne
Proces wysyłania zdarzeń z serwera do wyznaczonego serwera zdarzeń. Serwer zdarzeń kieruje zdarzenia do wyznaczonych odbiorników, takich jak programy użytkownika. Patrz także zdarzenie.

rekord zdarzenia
Rekord bazy danych opisujący rzeczywisty status i wynik zdarzenia.

rodzimy system plików
System plików lokalnie dodany do serwera plików, który nie został dodany do zarządzania pamięcią masową. Klient HSM nie udostępnia usług zarządzania pamięcią masową takiemu systemowi plików.

rozmyta kopia zapasowa
Wersja kopii zapasowej pliku, która może nie odzwierciedlać dokładnie bieżącej zawartości pliku, ponieważ kopia zapasowa tego pliku została utworzona w tym samym momencie, gdy plik był modyfikowany.

rozszerzanie
Zwiększanie ilość miejsca przeznaczonego do zapisania informacji w bazie danych lub informacji z dzienników odtwarzania.

S

SAN
Patrz sieć pamięci masowej.

Secure Sockets Layer (SSL)
Protokół zabezpieczeń zapewniający prywatność komunikacji. Używając protokołu SSL, aplikacje klient/serwer komunikują się w sposób chroniący przed podsłuchaniem, atakiem lub fałszowaniem komunikatów.

sekcja
Grupa wierszy w pliku, które spełniają tę samą funkcję lub definiują część systemu. Sekcje są zwykle oddzielone pustymi wierszami lub dwukropkami, a każda sekcja ma nazwę.

seria kopii zapasowych bazy danych
Jedna pełna kopia zapasowa bazy danych plus 32 przyrostowe kopie zapasowe utworzone od momentu utworzenia pełnej kopii zapasowej. Każda pełna kopia zapasowa rozpoczyna nową serię kopii zapasowych bazy danych. Każda seria kopii zapasowych jest identyfikowana przez liczbę. Patrz także obraz stanu bazy danych, pełna kopia zapasowa.

serwer
Program lub komputer udostępniający usługi innym programom lub komputerom. Patrz także klient.

serwer menedżera danych
Serwer gromadzący metadane dla spisu zasobów klienta i zarządzający transakcjami dla agenta pamięci masowej za pomocą sieci lokalnej. Serwer menedżera danych informuje agenta pamięci masowej o odpowiednich atrybutach biblioteki i identyfikatorze woluminu docelowego.

serwer nośników
W środowisku z/OS: program umożliwiający dostęp do dyskowej lub taśmowej pamięci masowej w systemie z/OS serwerem IBM Spectrum Protect działającym w systemach operacyjnych innych niż z/OS.

serwer plików
Dedykowany komputer wraz z peryferyjnymi urządzeniami pamięci masowej podłączonymi do sieci lokalnej, na których zapisywane są programy i pliki współużytkowane przez użytkowników sieci.

serwer plików NAS
Patrz serwer plików pamięci masowej podłączonej przez sieć.

serwer plików pamięci masowej podłączonej przez sieć (serwer plików NAS)

Dedykowane urządzenie pamięci masowej używające systemu operacyjnego zoptymalizowanego pod kątem funkcji udostępniania plików. Serwer plików NAS może jednocześnie posiadać cechy węzła i narzędzia przenoszenia danych.

serwer zarządzany
Serwer, który otrzymuje informacje konfiguracyjne z menedżera konfiguracji za pomocą subskrypcji jednego lub większej liczby profili. Informacje konfiguracyjne mogą obejmować definicje obiektów takich jak strategie i harmonogramy. Patrz także menedżer konfiguracji, konfiguracja korporacyjna, profil, subskrypcja.

serwer zdarzeń
Serwer, do którego inne serwery mogą wysyłać zdarzenia w celu ich zarejestrowania. Serwer zdarzeń kieruje zdarzenia do odbiorników skonfigurowanych dla zdarzeń serwera wysyłającego.

sesja
Logiczne lub wirtualne połączenie między dwoma stacjami, programami lub urządzeniami w sieci, które dwóm takim elementom umożliwia komunikację i wymianę danych. Patrz także sesja administracyjna.

sesja administracyjna
Okres, w którym użytkownik używający identyfikatora administratora komunikuje się z serwerem w celu wykonania czynności administracyjnych. Patrz także sesja węzła klienckiego, sesja.

sesja węzła klienckiego
Sesja, w której węzeł kliencki komunikuje się z serwerem w celu wykonania żądań utworzenia kopii zapasowej, odtwarzania, archiwizowania, pobierania, migrowania lub przywracania. Patrz także sesja administracyjna.

sieć lokalna (LAN)
Sieć łącząca szereg urządzeń w ograniczonym obszarze (takim jak pojedynczy budynek lub miasteczko akademickie), która może być połączona z rozleglejszą siecią.

sieć pamięci masowej (storage area network - SAN)
Dedykowana sieć pamięci masowej, która jest dostosowana do konkretnego środowiska i łączy serwery, systemy, produkty pamięci masowej, produkty sieciowe, oprogramowanie i usługi.

składnica danych
W środowisku zwirtualizowanym jest to położenie, w którym przechowywane są dane maszyny wirtualnej.

skrypt
Seria komend umieszczonych w pliku wykonujących konkretne funkcje podczas wykonywania tego pliku. Skrypty są interpretowane w czasie wykonywania. Patrz także skrypt komend produktu IBM Spectrum Protect.

skrypt komend produktu IBM Spectrum Protect
Sekwencja komend administracyjnych produktu IBM Spectrum Protect zapisana w bazie danych serwera IBM Spectrum Protect. Taki skrypt można uruchomić za pomocą dowolnego interfejsu umożliwiającego dostęp do serwera. Skrypt może zawierać podstawienia parametrów komend i instrukcje warunkowe. Patrz także plik makra, skrypt.

sprawdzanie poprawności
Sprawdzanie, czy w zestawie strategii nie występują warunki, które spowodują wystąpienie błędów, jeśli zestaw ten zostanie oznaczony jako aktywny. W procesie sprawdzania poprawności sprawdzane jest na przykład, czy zestaw strategii zawiera domyślną klasę zarządzania.

SSL
Patrz Secure Sockets Layer.

stabilizowany obszar plików
Obszar plików, który istnieje na serwerze, ale nie istnieje na kliencie.

stacja robocza
Terminal lub komputer osobisty, na którym użytkownik może uruchamiać aplikacje i który zwykle jest połączony z komputerem mainframe lub siecią.

stan globalnej nieaktywności
Stan wszystkich systemów plików, do których dodano zarządzanie pamięcią masową i w którym zarządzanie pamięcią masową zostało globalnie dezaktywowane dla węzła klienckiego.

stan pliku
Tryb zarządzania pamięcią masową pliku znajdującego się w systemie plików, do którego dodano zarządzanie pamięcią masową. Plik może być w jednym z trzech stanów: rezydentny, wstępnie zmigrowany lub zmigrowany. Patrz także migrowany plik, plik wstępnie zmigrowany, plik rezydentny.

stan systemu plików
Tryb zarządzania pamięcią masową systemu plików znajdującego się na stacji roboczej, w której zainstalowano klienta hierarchicznego zarządzania pamięcią masową (HSM). System plików może być w jednym z następujących stanów: rodzimy, aktywny, nieaktywny lub nieaktywny globalnie.

strona
Zdefiniowana jednostka miejsca na nośniku pamięci masowej lub w woluminie bazy danych.

subskrypcja
W środowisku pamięci masowej, proces identyfikowania subskrybentów, do których są dystrybuowane profile. Patrz także konfiguracja korporacyjna, serwer zarządzany.

surowy wolumin logiczny

Fragment woluminu fizycznego składający się z bloków nieprzydzielonych i nieposiadający definicji systemu plików JFS (Journaled File System). Wolumin logiczny jest dostępny do odczytu/zapisu jedynie przez funkcje we/wy niskiego poziomu. system plików podłączany automatycznie (AutoFS)

System plików zarządzany przez demona automounter. Demon automounter monitoruje określoną ścieżkę katalogów i automatycznie podłącza system plików w celu uzyskania dostępu do danych.

szeregowanie

Proces obsługi plików, które są modyfikowane podczas przetwarzania kopii zapasowych lub archiwalnych. Patrz także szeregowanie dynamicznych powtarzanych, szeregowanie statycznych powtarzanych, szeregowanie statyczne.

szeregowanie dynamiczne

Typ szeregowania kopii, w którym dla pliku lub folderu tworzona jest kopia zapasowa przy pierwszej próbie, bez względu na to, czy zostanie on zmieniony podczas wykonywania operacji tworzenia kopii zapasowej lub archiwalnej. Patrz także szeregowanie dynamicznych powtarzanych, szeregowanie statycznych powtarzanych, szeregowanie statyczne.

szeregowanie dynamicznych powtarzanych

Wartość szeregowania wskazująca, że nie może zostać utworzona kopia zapasowa ani archiwalna pliku, który został zmieniony podczas wykonywania operacji. Klient kopii zapasowych i archiwalnych ponawia operację tworzenia kopii zapasowej lub archiwalnej określoną liczbę razy. Jeśli plik zostanie zmieniony podczas każdej próby, klient kopii zapasowych i archiwalnych utworzy kopię zapasową lub archiwalną pliku dla ostatniej próby. Patrz także szeregowanie dynamiczne, szeregowanie, szeregowanie statycznych powtarzanych, szeregowanie statyczne.

szeregowanie statyczne

Wartość szeregowania grupy kopii określająca, że plik nie może być modyfikowany podczas wykonywania operacji tworzenia kopii zapasowej lub archiwalnej. Jeśli plik jest używany podczas pierwszej próby, menedżer pamięci masowej nie utworzy kopii zapasowej lub archiwalnej pliku. Patrz także szeregowanie dynamiczne, szeregowanie, szeregowanie dynamicznych powtarzanych, szeregowanie statycznych powtarzanych.

szeregowanie statycznych powtarzanych

Wartość szeregowania grupy kopii określająca, że plik nie może być modyfikowany podczas wykonywania operacji tworzenia kopii zapasowej lub archiwalnej. Klient podejmuje określoną liczbę prób wykonania tej operacji. Jeśli podczas każdej kolejnej próby plik jest używany, jego kopia zapasowa lub archiwalna nie zostanie utworzona. Patrz także szeregowanie dynamiczne, szeregowanie, szeregowanie dynamicznych powtarzanych, szeregowanie statyczne.

szybkie odtwarzanie VSS

Operacja, która odtwarza dane z lokalnego obrazu stanu. Obraz stanu jest kopią zapasową VSS, która znajduje się w lokalnym woluminie w tle. Operacja odtwarzania pobiera dane za pomocą metody kopiowania na poziomie plików.

szybkość przesyłania danych sieciowych

Szybkość obliczona przez podzielenie łącznej liczby przestanych bajtów przez czas ich przesyłania. Szybkość ta może być na przykład czasem przeznaczonym na przesyłanie danych w sieci.

szyfrowany system plików (Encrypted File System - EFS)

System plików, w którym używane jest szyfrowanie na poziomie systemu plików.

Ś

ścieżka

Obiekt definiujący relację jeden-do-jednego między miejscem źródłowym a miejscem docelowym. Ścieżka umożliwia dostęp do miejsca docelowego z miejsca źródłowego. Dane mogą być przesyłane z miejsca źródłowego do docelowego lub odwrotnie. Miejscem źródłowym może być na przykład narzędzie przenoszenia danych (takie jak serwer plików NAS), a miejscem docelowym może być napęd taśm.

T

TCA

Patrz zaufany agent komunikacyjny.

TCP/IP

Patrz protokół TCP/IP (Transmission Control Protocol/Internet Protocol).

tryb

Atrybut grupy kopii określający, czy ma być tworzona kopia zapasowa pliku, który nie uległ zmianie od chwili utworzenia ostatniej kopii zapasowej. Patrz także tryb bezwzględny, tryb zmodyfikowany.

tryb bezwzględny

W zagadnieniach związanych z zarządzaniem pamięcią masową: tryb grupy kopii zapasowej określający, że plik lub katalog jest uwzględniany w przyrostowej kopii zapasowej nawet wtedy, gdy nie został zmieniony od momentu utworzenia ostatniej kopii zapasowej. Patrz także tryb, tryb zmodyfikowany.

tryb częściowego przywracania plików

Tryb przywracania, w którym funkcja HSM odczytuje tylko fragment zmigrowanego pliku z pamięci masowej zgodnie z żądaniem aplikacji uzyskującej dostęp do pliku.

tryb dostępu

Atrybut puli pamięci masowej lub woluminu pamięci masowej określający, czy serwer może zapisywać dane w puli pamięci masowej lub woluminie pamięci masowej albo odczytywać je z tej puli lub woluminu.

tryb planowania

Typ operacji planowania dla węzła serwera i węzła klienckiego obsługujących dwa tryby planowania: z odpytywaniem przez klienta i inicjowane przez serwer.

tryb planowania inicjowany przez serwer

Technika komunikacji klient/serwer, w której serwer nawiązuje połączenie z węzłem klienckim, kiedy zachodzi konieczność wykonania zadań. Patrz także tryb planowania z odpytywaniem przez klienta.

tryb planowania z odpytywaniem przez klienta

Metoda działania, w której klient odpytuje serwer na potrzeby obsługi. Patrz także tryb planowania inicjowany przez serwer.

tryb zmodyfikowany

W zagadnieniach związanych z zarządzaniem pamięcią masową: tryb grupy kopii zapasowej określający, że plik lub katalog jest uwzględniany w przyrostowej kopii zapasowej tylko wtedy, gdy został zmieniony od momentu utworzenia ostatniej kopii zapasowej. Plik lub katalog jest uznawany za zmodyfikowany, jeśli uległa zmianie jego data, wielkość, właściciel lub uprawnienia. Patrz także tryb bezwzględny, tryb.

tworzenie kopii zapasowej odłączonego woluminu

Operacja tworzenia kopii zapasowej, w której wolumin jest zablokowany po to, aby żadne inne aplikacje w systemie nie mogły uzyskać do niego dostępu podczas wykonywania tej operacji.

tworzenie kopii zapasowej VSS

Operacja tworzenia kopii zapasowej używająca Usługi kopiowania woluminów w tle (VSS) firmy Microsoft. Operacja tworzenia kopii zapasowej tworzy obraz stanu (spójna kopia w danym momencie). Tę kopię można zapisać w lokalnych woluminach w tle lub w pamięci masowej serwera.

tworzenie kopii zapasowej woluminu podczas pracy

Operacja tworzenia kopii zapasowej, w której wolumin jest dostępny dla innych aplikacji w systemie podczas wykonywania tej operacji.

tworzenie kopii zapasowych z wykorzystaniem kroniki

Metoda tworzenia kopii zapasowych klientów Windows i AIX wykorzystująca mechanizm powiadamiania o zmianach w plikach w celu zwiększenia wydajności tworzenia przyrostowej kopii zapasowej przez ograniczenie konieczności pełnego skanowania systemu plików.

typ urządzenia plikowego

Typ urządzenia określający użycie plików o dostępie sekwencyjnym na dyskowej pamięci masowej jako woluminów.

U

UCS-2

Dwubajtowy (16-bitowy) schemat kodowania oparty na specyfikacji ISO/IEC 10646-1. UCS-2 definiuje trzy poziomy implementacji: poziom 1 - łączenie kodowanych elementów nie jest dozwolone; poziom 2 - łączenie kodowanych elementów jest dozwolone tylko dla następujących języków: tajski, indyjski, hebrajski i arabski; poziom 3 - dozwolona jest każda kombinacja kodowanych elementów.

UNC

Patrz uniwersalna konwencja nazewnictwa.

Unicode

Standard kodowania znaków umożliwiający wymianę, przetwarzanie i wyświetlanie tekstu napisanego w powszechnie stosowanych na świecie językach, a także w wielu tekstach klasycznych i historycznych.

unikalny identyfikator globalny (GUID)

Określony za pomocą algorytmu numer, który jednoznacznie identyfikuje jednostkę w systemie. Patrz także unikalny identyfikator uniwersalny.

unikalny identyfikator uniwersalny (UUID)

128-bitowy identyfikator liczbowy, dzięki któremu żadne dwa komponenty nie mają takiego samego identyfikatora. Patrz także unikalny identyfikator globalny.

uniwersalna konwencja nazewnictwa (UNC)

Nazwa serwera połączona z nazwą sieciową. Nazwy te razem identyfikują zasób w domenie.

uprawnienie

Prawo do dostępu do obiektów, zasobów lub funkcji. Patrz także klasa uprawnień.

usługa kroniki

W systemie Microsoft Windows: program śledzący zmiany w plikach znajdujących się w systemach plików.

usługa Volume Shadow Copy Service (VSS)

Zestaw funkcji API firmy Microsoft umożliwiających tworzenie kopii zapasowych w tle dla woluminów, dokładnych kopii zapasowych plików, w tym wszystkich plików otwartych itd.

UTF-8

Format UTF (Unicode Transformation Format) jest 8-bitowym formatem kodowania, który zaprojektowano z myślą o ułatwieniu użycia istniejących systemów wykorzystujących format ASCII. Wartością CCSID dla danych w formacie UTF-8

jest 1208.

utrata ważności

Proces identyfikowania plików, zestawów danych lub obiektów przeznaczonych do usunięcia z powodu upłynięcia ich daty ważności lub okresu przechowywania.

UUID

Patrz unikalny identyfikator uniwersalny.

uzgadnianie

Proces zapewniania spójności między oryginalnym repozytorium danych i większym systemem, w którym są przechowywane dane jako kopia zapasowa. Przykłady większych systemów, w których przechowywane są dane tworzące kopię zapasową, są serwery pamięci masowej lub inne systemy pamięci masowej. W procesie uzgadniania dane, które są określone jako zbędne, są usuwane.

użytkownik root

Użytkownik systemu wykonujący operacje bez ograniczeń. Użytkownik root ma specjalne prawa i uprawnienia wymagane do wykonania czynności administracyjnych.

V

VSS

Patrz usługa Volume Shadow Copy Service.

W

wartość procentowa dla migracji wstępnej

Ustawienie zarządzania pamięcią masową określające, czy kolejne zakwalifikowane do migrowania pliki w systemie plików zostaną poddane migracji wstępnej po przeprowadzeniu migracji progowej lub migracji na żądanie.

wersja

Kopia zapasowa pliku zapisana w pamięci masowej serwera. Najnowsza kopia zapasowa pliku jest wersją aktywną. Wcześniej utworzone kopie tego samego pliku są wersjami nieaktywnymi. Liczba wersji przechowywanych przez serwer zależy od atrybutów grupy kopii określonych w klasie zarządzania.

wersja aktywna

Najnowsza zapisana kopia zapasowa pliku. Aktywnej wersji pliku nie można usunąć aż do momentu, gdy proces tworzenia kopii zapasowej wykryje, że użytkownik zastąpił ten plik nowszą wersją lub usunął go z serwera plików albo stacji roboczej. Patrz także wersja kopii zapasowej, wersja nieaktywna.

wersja kopii zapasowej

Plik lub katalog, dla którego węzeł kliencki utworzył kopię zapasową w pamięci masowej. W pamięci masowej może istnieć więcej niż jedna wersja kopii zapasowej, ale tylko jedna z nich jest wersją aktywną. Patrz także wersja aktywna, grupa kopii, wersja nieaktywna.

wersja nieaktywna

Wersja kopii zapasowej pliku, która nie jest najnowszą wersją kopii zapasowej albo jest wersją kopii zapasowej pliku, która już nie istnieje w systemie klienta. Nieaktywne wersje kopii zapasowych są kwalifikowane do przetwarzania obsługującego utratę ważności zgodnie z klasą zarządzania przypisaną do danego pliku. Patrz także wersja aktywna, wersja kopii zapasowej.

węzeł

Serwer plików lub stacja robocza zarejestrowana na serwerze, na której zainstalowano program klienta kopii zapasowych i archiwalnych.

węzeł agenta

Węzeł kliencki, któremu nadano uprawnienie proxy do wykonywania operacji w imieniu innego węzła klienckiego będącego węzłem docelowym.

węzeł docelowy

Węzeł kliencki, do którego innym węzłem klienckim (zwanym węzłami agentów) nadano uprawnienie proxy. Uprawnienie proxy umożliwia węzłom agentów wykonywanie operacji, takich jak tworzenie i odtwarzanie kopii zapasowych, w imieniu węzła docelowego będącego właścicielem danych.

węzeł kliencki

Serwer plików lub stacja robocza zarejestrowana na serwerze, na której zainstalowano program klienta kopii zapasowych i archiwalnych.

węzeł NAS

Węzeł kliencki, który jest serwerem plików NAS. Dane węzła NAS są przesyłane przez serwer plików NAS sterowany przez protokół zarządzania danymi sieciowymi (NDMP). Węzeł NAS jest zwany również węzłem serwera plików NAS.

węzeł serwera plików NAS

Patrz węzeł NAS.

wiek pliku

W kontekście ustalania priorytetów migracji: liczba dni od ostatniego dostępu do pliku.

wielkość pliku pośredniczącego

Wielkość pliku, który zastępuje plik oryginalny w lokalnym systemie plików, gdy zostanie on zmigrowany do pamięci masowej serwera. Wielkość zdefiniowana dla pliku pośredniczącego określa ilość danych wiodących, które można zapisać w pliku pośredniczącym. Domyślnym ustawieniem wielkości pliku pośredniczącego jest wielkość bloku zdefiniowana dla systemu plików pomniejszona o 1 bajt.

wielkość szacunkowa

Wyrażona w megabajtach ilość miejsca dostępnego w puli pamięci masowej.

wirtualny obszar plików

Reprezentacja katalogu w systemie plików NAS w postaci ścieżki do tego katalogu.

wirtualny punkt podłączenia

Gałąź katalogu systemu plików, która jest zdefiniowana jako wirtualny system plików. Wirtualny system plików jest obsługiwany przez własny obszar plików na serwerze. Serwer przetwarza wirtualny punkt podłączenia jako oddzielny system plików, ale system operacyjny klienta tego nie robi.

wirtualny system plików LOFS (loopback virtual file system - LOFS)

System plików tworzony przez podłączenie katalogu do innego katalogu lokalnego; zwany również podłączeniem przez podłączenie. System plików LOFS można także utworzyć za pomocą programu automounter.

wolumin

Indywidualna jednostka pamięci masowej na dysku taśmie lub innym nośniku, na którym możliwe jest zapisywanie danych, obsługująca określoną postać listy identyfikatorów i parametrów, takiej jak etykieta woluminu lub sterowanie we/wy. Patrz także wolumin chwilowy, pamięć masowa serwera, pula pamięci masowej, wolumin puli pamięci masowej.

wolumin chwilowy

Wolumin oznaczony etykietą, który jest pusty lub nie zawiera żadnych poprawnych danych, nie jest zdefiniowany i jest dostępny do użycia. Patrz także wolumin.

wolumin logiczny

Część woluminu fizycznego zawierająca system plików.

wolumin puli pamięci masowej

Wolumin przypisany do puli pamięci masowej. Patrz także pula danych aktywnych, pula pamięci masowej kopii, podstawowa pula pamięci masowej, pamięć masowa serwera, wolumin.

wolumin w tle

Dane zapisane z obrazu stanu woluminu. Obraz stanu może być tworzony również wtedy, gdy aplikacje w systemie zapisują dane w woluminie.

wolumin wirtualny

Plik archiwum na serwerze docelowym, który reprezentuje wolumin nośników sekwencyjnych dla serwera źródłowego.

WPAR

Patrz partycja zarządzania obciążeniem.

wtyczka

Osobno instalowany moduł oprogramowania, który dodaje funkcję do istniejącego programu, aplikacji lub interfejsu.

WWN

Patrz nazwa WWN.

wybiórcze tworzenie kopii zapasowych

Proces tworzenia kopii zapasowych niektórych plików lub katalogów z domeny klienta. Tworzona jest kopia zapasowa tych plików, które nie zostały wykluczone na liście włączeń/wykluczeń. Pliki te muszą spełniać warunki dotyczące szeregowania w grupie kopii zapasowych klasy zarządzania przypisanej do każdego pliku. Patrz także przyrostowa kopia zapasowa.

wykluczanie

Proces identyfikowania plików na liście włączeń/wykluczeń. Ten proces zapobiega tworzeniu kopii zapasowych lub migrowaniu plików za każdym razem, gdy użytkownik wprowadzi operację przyrostowego lub selektywnego tworzenia kopii zapasowych. Plik można wykluczyć z operacji tworzenia kopii zapasowej i z zarządzania pamięcią masową, tylko z operacji tworzenia kopii zapasowej lub tylko z zarządzania pamięcią masową.

Z

zadanie migracji

Specyfikacja plików do zmigrowania i działań, które mają być wykonane na plikach macierzystych po migracji. Patrz także plik zadania, migracja progowa.

zagregowana szybkość przesyłania danych

Statystyka wydajności wskazująca średnią liczbę bajtów przestanych w ciągu sekundy podczas przetwarzania danej operacji.

zapis lustrzany

Proces zapisywania tych samych danych w wielu dyskach jednocześnie. Zapis lustrzany chroni przed utratą danych uwzględnionych w dzienniku odtwarzania lub bazie danych.

zarządzanie cyklem życia informacji (information lifecycle management - ILM)

System zarządzania plikami wykorzystujący strategie dla pul pamięci masowej i zestawów plików. Patrz także General Parallel File System.

zarządzanie pamięcią masową
Patrz hierarchiczne zarządzanie pamięcią masową.

zaufany agent komunikacyjny (TCA)
Program obsługujący protokół hasła logowania, gdy klienci używają funkcji generowania hasła.

zdalne
W produktach HSM: termin odnoszący się do miejsca źródłowego migrowanych plików, które są przenoszone. Patrz także lokalne.

zdarzenie
Wystąpienie czegoś istotnego dla zadania lub systemu. Zdarzenia mogą obejmować zakończenie lub niepowodzenie operacji, działanie użytkownika lub zmianę stanu procesu. Patrz także rejestrowanie korporacyjne, odbiornik.

zestaw kopii zapasowych
Przenośna, skonsolidowana grupa aktywnych wersji plików kopii zapasowej wygenerowanych dla klienta kopii zapasowych i archiwalnych.

zestaw opcji klienta
Grupa opcji zdefiniowanych na serwerze i używana w węzłach klienckich razem z plikami opcji klientów.

zestaw strategii
Grupa reguł w domenie strategii. Reguły te określają sposób automatycznego zarządzania danymi lub zasobami pamięci masowej dla węzłów klienckich w domenie strategii. Reguły mogą być zawarte w klasach zarządzania. Patrz także aktywny zestaw strategii, klasa zarządzania.

zestaw węzłów GPFS
Podłączona, zdefiniowana grupa systemów plików GPFS

znak wieloznaczny
Znak specjalny, np. gwiazdka (*) lub znak zapytania (?), który może reprezentować jeden lub wiele znaków. Znak wieloznaczny może być zastąpiony dowolnym znakiem lub zestawem znaków.

znak zastępczy
Patrz znak wieloznaczny.

zużycie zasobów sesji
Czas oczekiwania, czas procesora i ilość pamięci zużyte lub pobrane podczas sesji klienta.